WEAVE SPRINT
PAPER IV

MD and DNB psychiatry, theory revision

Neurology, Medicine and Recent Advances

Where psychiatry meets the rest of medicine: epilepsy and the dementias, the endocrine and systemic diseases that reach the mind, the drugs and what to do when they fail, and psychiatry at the scale of a population. Eleven complete sets of notes, in one book.

Dr. Wilfred D'souza · Dr. Niharika Reddy

WEAVE
CENTRE FOR INTEGRATIVE PSYCHIATRY

Weave
Centre for Integrative Psychiatry

WEAVE SPRINT
PSYCHIATRY OUTREACH PROGRAMME

Psychiatry that listens before it prescribes.

Therapy-informed psychiatry for adults, children, couples and families. Online across India, in person in Mumbai.





Assessments

ADHD, autism and developmental. Anxiety, mood, attention and personality patterns.



Consultations and sessions

Psychiatric consultations, therapy sessions, medication review and couples sessions.



Family Library

Free plain-language guides for families, in five languages.

`weave.clinic/family`

REFER, OR SEND SOMEONE OUR WAY

Your patient contacts Weave directly and decides whether to book.

You remain the referring clinician. Any coordination between clinicians happens with the patient's permission.

CONTACT

`weave.clinic/contact`
`WhatsApp +91 7625004229`

Dr. Niharika Reddy and Dr. Wilfred D'souza, Consultant Psychiatrists. Runwal Forests Clinic, Kanjurmarg West, Mumbai, and online.

WEAVE.CLINIC

Routine referrals only. This is not an emergency route.

WEAVE SPRINT · PAPER IV

Neurology, Medicine and Recent Advances

THE WEAVE SPRINT SERIES

TITLE

Weave Sprint, Paper IV: Neurology, Medicine and Recent Advances

Part of the Weave Sprint series, MD and DNB psychiatry theory revision.

WRITTEN BY

Dr. Wilfred D'souza and Dr. Niharika Reddy

PUBLISHED BY

Weave, Centre for Integrative Psychiatry · weave.clinic

EDITION

First compiled 2026. Free to read, share and print for study.

The clinical content is written for doctors and doctors in training. It is a study and revision aid, not a substitute for current guidelines or clinical judgement. Verify doses and thresholds against the current references and your local protocols before any clinical use.

For every resident who has ever sat up the night before,
afraid it will not be enough.

ACKNOWLEDGEMENT

A word of thanks

I want to begin by thanking something I cannot quite name. Call it whatever you like. I have never been certain what to call it, and I have stopped needing to be. There is a part of this work that does not feel like mine. The line that arrives before I have thought of it, the connection I did not plan, the steadiness on a morning I had no reason to be steady. I only hold the pen. Something moves through the holding, and I have learned to thank it without asking it for a name.

Someone once asked me why I give all of this away. Why hand out the notes I sat up making, to the very people sitting the same exam beside me. It is a fair question. If I had kept them, I would probably have scored a few marks more than the next person, and I might have been thought of as the better student. I have turned that over honestly, and I have decided it is a small thing to want. What I get instead, by putting the work in your hands, is the chance to take a little of the fear out of a great many nights. That trade is not close. Easing that fear is not a detour on the way to becoming a good psychiatrist. It is the thing itself. I did not come here to be the best student in the room. I came here to be a good psychiatrist.

And to Niharika, who reads everything first and knows when to be honest and when to wait; and to the residents who study from these pages and write back with the corrections and the questions: thank you. The book is better for you, and so am I.

Dr. Wilfred D'souza

Foreword

This is the fourth paper of the written examination, gathered into one book. The first paper laid the ground the discipline stands on, the second was the discipline at work on the disorders themselves, and the third was psychiatry practised at its edges. This one is psychiatry meeting the rest of medicine, and it is the paper where the question most often is not which psychiatric disorder this is but whether it is a psychiatric disorder at all: the seizure that presents as a psychosis, the gland driving the mood, the failing organ that takes the cognition with it, the drug producing the syndrome it was given to treat. Half of this paper is about drugs, which is fitting, because prescribing is where the most confident errors in psychiatry are made. The rest is about noticing, in a picture a psychiatric formulation would absorb without complaint, the thing that does not belong to psychiatry at all.

The notes were written to be understood, not just revised. Each chapter is a full account of its subject, from the mechanism up to the point the exam tests, so that you can read to learn the first time and read to recall every time after. Where a fact is worth holding exactly, it is marked. Where a mistake is easy to make, it is named before you make it. The colour code on the next pages tells your eye what kind of fact it is looking at before your mind has to decide, and after a few pages it stops being a key you consult and becomes a sense you carry.

Read it in order if you are meeting the material, or open it anywhere if you are returning to it. The contents at the front and the index at the back both jump to the page you want. However you use it, the hope held in these pages is plain: that one ordered, trustworthy place to study from makes the work a little lighter, and that you walk into the hall steadier than you would have otherwise.

HOW TO USE THIS BOOK

The colour memory code

Every note in this volume speaks one visual language, so once you learn it in the first chapter you read the rest without thinking about it.

Concept and rule	The core idea and the principle that follows from it. When a line is petrol, it is the thing to understand, and the thing worth being able to state.
Key number	A figure to hold exactly: a dose, a threshold, a prevalence, a cut-off. Amber means memorise this number, not just the shape of it.
Trap	The mistake the exam is waiting for. When a line is coral, it is warning you off an error that looks right until it costs you the mark.

Hold this is the single line from a topic you must carry out of it. **Good to remember** gathers, for each named thing, what it is, when it matters, and the one number or formula worth keeping. Both help a tired brain find the centre of a dense page fast.

The contents at the front is clickable, and so is the index at the back: every entry jumps to its page. Read to learn once, and to recall as often as you need.

Contents

FRONT MATTER

Foreword 7

How to use this book 8

PART I **Epilepsy and Psychiatry** 14

- Epilepsy for the psychiatrist: classification, localisation and the interface 14
- The peri-ictal axis: syndromes by their timing to the seizure 63
- The psychoses of epilepsy 111
- Mood, anxiety and suicide in epilepsy 142
- Personality, behaviour and cognition in epilepsy 171
- Psychogenic non-epileptic seizures 207
- Antiepileptic drugs: the psychiatric interface 242
- Consolidate and apply 272
- References 285

PART II **Neurology Foundations for Psychiatrists** 287

- The bedside examination, what it localises, and the investigations it earns 287
- Headache and the psychiatric interface 341
- Sleep: the architecture and its disorders 368
- Consolidate and apply 395
- References 408

PART III **Dementias and Neurocognitive Disorders** 411

- The major dementias, from Alzheimer's to the frontotemporal syndromes 411
- The treatable, the rapid and the borderline 509
- Assessment, biomarkers and the pharmacology of dementia 553

Consolidate and apply 611

References 626

PART IV **Endocrine and Metabolic Psychiatry** 628

The endocrine axes: adrenal, thyroid and parathyroid 628

Metabolic and nutritional disease 688

The metabolic cost of treatment, diabetes and electrolytes 733

Encephalopathies, reproductive endocrinology and toxins 787

Consolidate and apply 828

PART V **Stroke, TBI, Delirium and Focal Brain Syndromes** 844

The cerebrovascular and post-stroke syndromes 844

Traumatic brain injury and its psychiatric sequelae 887

Delirium and the amnestic syndromes 929

The bedside examination and the focal lobe syndromes 985

Structural lesions and the seizure overlap 1029

Consolidate and apply 1051

Sources and further reading 1064

PART VI **Neuropsychiatry of Systemic and Neurological Disease** 1067

HIV and its neuropsychiatric syndromes 1067

Parkinson disease and its neuropsychiatry 1104

Huntington disease and other movement disorders 1134

Autoimmune and infective encephalitides 1157

Neurosyphilis and CNS infections 1183

Multiple sclerosis and psychoneuroimmunology 1201

Degenerative, neuromuscular and functional neurology 1213

Consolidate and apply 1237

Sources and further reading 1250

PART VII **Psychopharmacology I: Principles and Core Drug Classes** 1252

Receptor pharmacology and pharmacokinetic principles 1252
Antipsychotics 1288
Antidepressants 1334
Mood stabilisers 1367
Anxiolytics, hypnotics and other agents 1399
Consolidate and apply 1427
Sources and further reading 1439

PART VIII **Psychopharmacology II: Emergencies, Resistance, Monitoring and Interactions** 1442

Drug emergencies 1442
Treatment resistance and augmentation 1478
Drug monitoring, pharmacovigilance and pharmacogenomics 1508
Interactions and prescribing in physical illness 1531
Neurostimulation and device-based treatments 1570
Pharmacogenetic and structural drug-reaction syndromes 1609
Consolidate and apply 1624

PART IX **Recent Advances and Landmark Studies** 1647

Recent advances and landmark studies 1647
Individual newer agents 1747
Consolidate and apply 1783

PART X **Community Psychiatry and Public Health** 1797

National programmes and mental health legislation 1797
Epidemiology, burden and treatment gap 1834
Suicide prevention and farmer's suicide 1852
Stigma and anti-stigma interventions 1876
Psychiatric rehabilitation and community care 1893
Primary care, integration and service delivery models 1921
Special-setting and disaster mental health 1945

Prevention, economics and digital mental health 1969
Culture, global mental health and traditional healing 1996
Consolidate and apply 2007

PART XI **Medical and Systemic-Disease Psychiatry** 2022

Cardiac, respiratory and critical-illness psychiatry 2022
Renal, transplant and perioperative psychiatry 2056
Oncology, palliative and paraneoplastic psychiatry 2069
Autoimmune and inflammatory systemic disease, and burn injury 2083
Consolidate and apply 2111

BACK MATTER

Subject index 2127
About the authors 2145

PART I

Epilepsy and Psychiatry

The two-way traffic between seizures and the mind: the peri-ictal syndromes, the psychoses of epilepsy, non-epileptic attacks, and prescribing across both illnesses at once.

01 Epilepsy for the psychiatrist: classification, localisation and the interface

02 The peri-ictal axis: syndromes by their timing to the seizure

03 The psychoses of epilepsy

04 Mood, anxiety and suicide in epilepsy

05 Personality, behaviour and cognition in epilepsy

06 Psychogenic non-epileptic seizures

07 Antiepileptic drugs: the psychiatric interface

08 Consolidate and apply

09 References

VISUAL 1.1



The three possible authors of a psychiatric symptom

- Seizure-linked
- Treatment-linked
- Independent psychiatric condition
- Timing and collateral
- Reassess

A psychiatric symptom in epilepsy may arise from the seizure disorder, its treatment, or an independent condition; timing and collateral evidence determine attribution.

Source lock: day-30.final.html#sec-t01

VISUAL 1.2



The peri-ictal axis

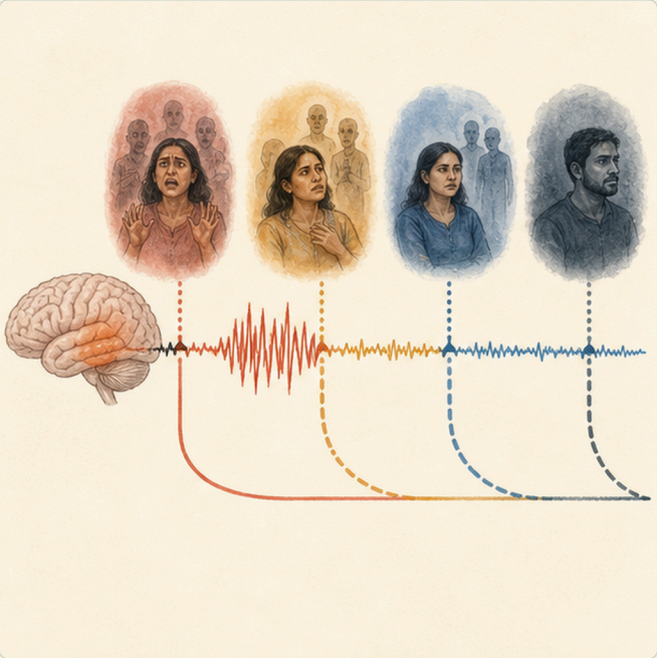
- Preictal
- Ictal
- Postictal
- Interictal
- Lucid interval

Peri-ictal psychiatric syndromes are classified by their temporal relation to the seizure; the lucid interval is especially important in postictal psychosis.

Source lock: day-30.final.html#sec-t09

VISUAL 1.3

Psychosis across the seizure timeline



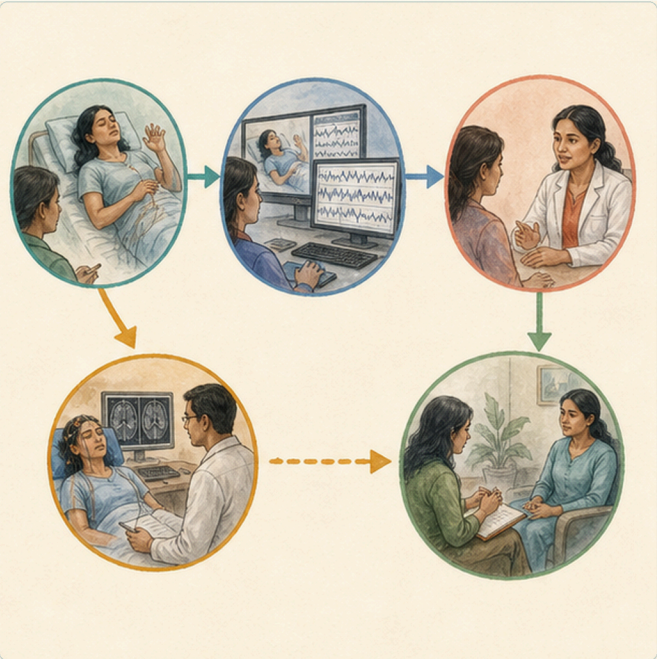
- Ictal
- Postictal
- Interictal
- Primary psychosis
- Seizure timeline

Psychosis associated with epilepsy is interpreted through timing, seizure relationship, recovery pattern and longitudinal course rather than symptom form alone.

Source lock: day-30.final.html#sec-t19

VISUAL 1.4

PNES is a positive diagnosis



- Positive semiology
- Video-EEG
- Communicate the diagnosis
- Psychological treatment
- Coexisting epilepsy

PNES is diagnosed positively, communicated clearly and treated psychologically while coexisting epilepsy remains an explicit possibility.

Source lock: day-30.final.html#sec-t37

PAPER IV

Epilepsy and psychiatry

Seven bands . one complete notes document

Epilepsy for the psychiatrist: classification, localisation and the interface

1 · Epilepsy and psychiatry: the interface, and how to read this chapter

Epilepsy is the disorder neurology and psychiatry have never managed to divide cleanly between them. A seizure is an electrical event, but the cortex it discharges through is the same organ that generates mood, thought, perception and personality, so a large part of what epilepsy does to a patient lands squarely in the psychiatrist's clinic. Candidates lose marks in this territory not because the facts are hard but because the material sits astride a border they were trained to treat as a wall. This chapter is written to dissolve that wall. It is a standalone reference: it assumes you can already recognise a seizure, and asks instead what the psychiatrist must add once epilepsy is on the table.

The examiner is rarely testing whether you can list the psychiatric complications of epilepsy, because any textbook can supply that list. The real test is whether you can take a single presenting symptom - a wave of fear, a hallucination, a spell of low mood, an aggressive outburst - and reason about it the way a clinician who understands epilepsy would: in time, in space, and by cause. Those three questions organise everything that follows, and this opening section exists to install them before the detail arrives. A reader who reaches the later sections without them will be forced to memorise each syndrome as an isolated fact; a reader who carries them will find the same syndromes deducible, which is both faster and more durable under examination pressure.

1.1 Why this material sits in the psychiatry syllabus

For most of recorded medical history epilepsy was construed as a disorder of the soul long before it was understood as a disorder of the cortex, and the **interface** between seizures and psychiatric phenomena is therefore not a modern administrative overlap but one of the oldest observations in clinical medicine. The reason it endures is structural rather than historical. The circuits that discharge in the commonest human epilepsies are the same limbic and paralimbic systems that regulate emotion, memory and drive, so disturbance in those circuits announces itself in the currency of psychiatry: fear, elation, depersonalisation, psychosis, altered sexuality and changed personality. A psychiatric symptom in a patient with epilepsy is frequently not a coincidental second illness but a direct expression of the seizure disorder or of the brain that hosts it.

This matters for the examination because it reframes the whole subject. Psychiatric phenomena in epilepsy are best approached not as a catalogue of comorbidities to be memorised but as the predictable behavioural output of a known lesion in a known circuit. Once that stance is adopted, the individual syndromes stop being arbitrary and start being derivable. The chapter is organised to reward exactly that reasoning, and the sections ahead repeatedly ask you to move from the

seizure to the symptom and back, rather than treating the two as separate examinable objects that happen to share a patient. It is also why the psychiatrist cannot delegate this material to the neurologist and vice versa: the seizure and the symptom are the same event described in two vocabularies, and the patient is served only when someone reads both.

1.2 A relationship that runs in both directions

The single most important conceptual correction this chapter makes is that the link between epilepsy and psychiatric disorder is not a one-way causal arrow but a genuinely **bidirectional relationship**. The obvious direction is familiar: living with a seizure disorder raises the risk of depression, anxiety, psychosis and suicide, through mechanisms that are partly neurobiological and partly the burden of a chronic, stigmatised, unpredictable illness. The direction candidates neglect is the reverse one. A history of a primary psychiatric disorder, depression in particular, precedes and predicts the later onset of epilepsy, which cannot be explained by the psychological reaction to seizures because the mood disorder comes first. The two conditions appear to share vulnerability in the same circuits, so each raises the probability of the other.

Holding both arrows in mind changes clinical practice. It means a psychiatric history is part of the epilepsy assessment and a seizure history is part of the psychiatric assessment, and it means the two disorders are best thought of as neighbours on a shared substrate rather than as cause and consequence. The epidemiological weight of this two-way traffic - how much commoner these disorders actually are in epilepsy, and how the risk is distributed - is quantified where the chapter turns to prevalence, incidence and the bidirectional data; this section installs the frame, and that section supplies its magnitude.

IN INDIA

The bidirectional interface is not an abstraction in Indian practice. Epilepsy still carries heavy social stigma that bears on marriage, employment and schooling, and much of that burden is psychological and therefore the psychiatrist's to manage. Care is frequently split across a neurologist or general physician and a psychiatrist with little formal liaison, and the psychiatrist is often the first clinician a patient reaches when the presentation is behavioural rather than convulsive. Reading the relationship as two-way, rather than assuming the psychiatric picture is merely a reaction to fits, is what keeps that first contact from missing a treatable seizure disorder.

1.3 The first question - when, relative to the seizure?

The organising spine of the entire subject is temporal. Almost every psychiatric phenomenon in epilepsy is defined less by its cross-sectional appearance than by its timing against the seizure, and the chapter therefore introduces a **peri-ictal temporal axis** as its master framework. The axis sorts any symptom into one of four windows according to when it occurs relative to the fit: before it, during it, immediately after it, or in the seizure-free intervals between attacks. A fear that is **ictal** is the seizure itself; the same fear arising in the settled periods between seizures is an **interictal** anxiety phenomenon with a different mechanism, course and treatment. The words are similar; the disorders are not.

Pre	Ictal	Post	Inter
BEFORE THE SEIZURE - PRODROME AND PREICTAL CHANGE	DURING THE SEIZURE - THE DISCHARGE ITSELF	AFTER THE SEIZURE - THE RECOVERY WINDOW	BETWEEN SEIZURES - THE SETTLED INTERVAL

This axis is the framework that the peri-ictal band of the chapter unfolds in full, placing each syndrome - the preictal prodrome, the ictal affective and psychotic phenomena, postictal confusion and psychosis, and the interictal disorders - into its correct window. For orientation, only the discipline of the question matters: get into the habit of establishing the temporal relationship of a symptom to the last seizure before you reach for a diagnostic label, because the label often changes entirely once the timing is known.

RULE **Timing comes before naming.** A psychiatric symptom in a patient with epilepsy is characterised first by where it sits on the peri-ictal axis and only then by its content. The same phenomenology means different things in different windows, so the temporal relationship to the seizure is the first datum to establish, not an afterthought once the diagnosis has been guessed.

1.4 The second question - where is the focus?

Once timing is settled, location does much of the remaining explanatory work. Epilepsy is not a single disease but a family of disorders, and the psychiatric colour of a given case depends heavily on where in the brain the seizures begin. The chapter opens with classification precisely because the distinction between focal onset and generalised onset, and within focal epilepsy the identity of the lobe, is what predicts the psychiatric phenotype. **Localisation** is therefore not a neurologist's technicality that psychiatrists can skip; it is the variable that tells you which behavioural syndrome to expect.

Two regions carry most of the psychiatric salience, and each has its own section. Temporal lobe epilepsy, arising from limbic and mesial temporal structures, is the substrate for the richest affective, psychotic and personality phenomena, while frontal lobe epilepsy produces its own striking behavioural and motor presentations that are easily mistaken for functional or primary psychiatric disorder. The chapter also devotes a mechanism section to the limbic system, kindling and the neurobiology that links repetitive seizure activity to durable behavioural change; the full structural anatomy of that system belongs to the Functional Neuroanatomy and Neurophysiology notes and is cross-referenced there rather than rebuilt here. Reading the **semiology** of a seizure as a map to its origin is the skill these sections train.

The practical lesson of localisation is that the psychiatrist should never receive the word "epilepsy" as a sufficient diagnosis. Which epilepsy, from where, matters as much as the fact of seizures, because a mesial temporal focus and a frontal focus generate different behavioural risks and warrant different anticipatory questions. When the referring information gives only the bare label, the localising detail - the aura, the ictal behaviour, the imaging and the electroencephalographic focus - is worth actively seeking, because it is frequently the single fact that makes the psychiatric presentation intelligible.

1.5 The third question - disease, drug or neither?

The final master question is the one that separates safe clinicians from dangerous ones, and it recurs so often that the chapter is deliberately built around it. Any psychiatric symptom in a patient with epilepsy has three broad possible authors, and **attribution** to the correct one dictates management. The symptom may be produced by the epilepsy itself, whether as a direct peri-ictal event or as an interictal consequence of the seizure disorder. It may be produced by the treatment, since the antiepileptic drugs carry their own well-described psychiatric and cognitive adverse effects. Or the paroxysmal event may not be epileptic at all, but a functional, non-epileptic attack that mimics a seizure.

This attribution problem is why several parts of the chapter are built as explicit, disjoint pairs: cognitive change from the disease set against cognitive change from the drugs; mood and suicidality intrinsic to epilepsy set against mood and suicidality attributable to antiepileptic medication. In each case the examiner is testing whether you can assign the same clinical picture to disease or to drug, because the two demand opposite responses. The point of naming the three authors at the outset is to make you ask, of every symptom, which one you are looking at before you act on it.

TRAP **Blaming the disease for what the drug is doing, or the reverse.** Candidates confidently ascribe apathy, slowed cognition or low mood to "the epilepsy" when the antiepileptic drug is the cause, or reassure that a symptom is a harmless drug effect when it is the seizure disorder declaring itself. The cross-sectional picture cannot settle it; only a careful account of onset, dose changes and timing can.

PITFALL

Adjusting psychotropic medication in a patient with epilepsy without speaking to the treating neurologist. A change made in isolation can destabilise seizure control, compound an antiepileptic drug's adverse effect, or collide with an enzyme interaction that the psychiatrist never saw coming. The interface is bidirectional in management too: neither clinician should prescribe across it blind to the other.

1.6 A map of the chapter

With the three questions in place, the architecture of what follows becomes legible. The chapter moves in bands, each answering the master questions for a defined slice of the material, and the accompanying map sets out how they fit together. The early bands establish the substrate - how seizures are classified and localised and how the two conditions relate epidemiologically. The middle bands work through the psychiatric syndromes in the order the peri-ictal axis dictates, then the psychoses, then mood, anxiety and suicide, then the personality, behavioural and cognitive changes. The later bands address the non-epileptic attacks that mimic all of the above, and finally the antiepileptic drugs viewed as objects of psychiatric interest in their own right.

BAND OF THE CHAPTER	WHAT IT COVERS	THE DISCRIMINATION IT DRILLS
Classification and localisation	how seizures are classified and where in the brain they arise	focal versus generalised onset; which lobe is involved
The peri-ictal axis	psychiatric syndromes sorted by their timing to the seizure	ictal versus postictal versus interictal
The psychoses of epilepsy	ictal, postictal and chronic interictal psychotic states	epileptic psychosis versus primary schizophrenia
Mood, anxiety and suicide	depression, anxiety disorders and suicide risk in epilepsy	intrinsic to the disease versus attributable to the drug
Personality, behaviour and cognition	interictal personality change, aggression and cognitive effects	the epilepsy versus its treatment versus surgery
Non-epileptic seizures	functional, psychogenic attacks and their management	epileptic versus non-epileptic events
Antiepileptic drugs at the interface	the drugs as mood agents, as offenders and as interacting molecules	therapeutic effect versus adverse psychiatric effect

The chapter closes with worked vignettes and an active-recall consolidation section, so the reading load is front-weighted and the application is deliberately saved for last. Nothing in the map needs memorising as a list; it is here so that when a later section names a syndrome you already know which question it is answering and which neighbour it borders.

1.7 What the examination actually rewards

The subject has a small number of recurring intellectual moves that earn most of the marks, and recognising them in advance is worth more than any single fact. Across every band, the questions that discriminate strong candidates from weak ones ask essentially the same things in different clothing, so it is worth naming them plainly before the detail begins.

- Placing a symptom on the peri-ictal axis before diagnosing it, so that an ictal, postictal and interictal version of the same complaint are never confused.
- Reading localisation and semiology together to predict the psychiatric phenotype, rather than treating the lobe as a neurological detail.
- Attributing a symptom to the disease, the drug or a functional cause, since the same picture demands opposite management under each.
- Discriminating the epileptic from the primary psychiatric disorder that it mimics, using course and relationship to seizures rather than cross-sectional appearance.
- Coordinating psychiatric treatment with seizure control, so that neither is optimised at the other's expense.

A candidate who carries these five habits into the rest of the chapter will find that the individual syndromes fall into place as instances of a general method, and that the awkward look-alikes the examiner favours become tractable rather than treacherous. The detail still has to be learned, but it is learned faster when it hangs on a frame.

GOOD TO REMEMBER

Take a seizure history from every patient who presents with episodic or paroxysmal psychiatric symptoms, and in every patient with known epilepsy establish the temporal relationship of any new psychiatric symptom to the last seizure before you name it. Those two habits, applied on Monday morning, catch most of the errors this chapter is written to prevent.

1.8 Holding it together

Everything in this chapter can be reduced to a single disciplined sequence applied to whatever symptom the patient brings. First locate the symptom in time against the seizure, then locate the seizure in space within the brain, then decide whether the disease, the drug or a functional process is its author. Only once those three are settled does the diagnostic label become safe to attach, because each of them can change the label entirely. The syndromes that fill the rest of the chapter are simply what this sequence produces when it is run on the real presentations epilepsy generates. None of the sequence is exotic; its power lies in being applied consistently, to the fear and the hallucination and the outburst alike, so that no presentation is ever named on its surface appearance before its timing, its origin and its cause have been established.

MNEMONIC

Localise, Time, Attribute

The three master questions of the whole chapter, in the order the clinic uses them. **Localise**: where do the seizures arise, and what phenotype does that predict? **Time**: where on the peri-ictal axis does this symptom sit? **Attribute**: is its author the disease, the drug, or a non-epileptic process? Run all three before you name anything, and the look-alikes that trap candidates come apart in your hands.

Before you name a psychiatric symptom in a patient with epilepsy, place it in time against the seizure and in space against the focus, and decide whether disease, drug or a functional process produced it.

2 · ILAE seizure classification (focal, generalised, unknown onset) and clinical relevance to psychiatry

Classification is where the clinical thinking begins, not where it ends. When the International League Against Epilepsy last recast how seizures are named, in 2017, it fixed the discipline's opening question as something deceptively plain: where does the seizure start? Every seizure is now sorted by its **point of onset** into one of three families, focal onset, generalized onset, or unknown onset, and that single decision organises almost everything downstream of it. It shapes the syndrome the patient is likely to have, the causes worth pursuing, the investigations that earn their keep, and the treatment that is rational rather than reflexive. For the psychiatrist this is not neurology's private bookkeeping. A seizure and a psychiatric episode can present as near-identical events across the same consulting-room desk, and the only disciplined way to tell them apart is to know what a seizure is, how it is classified, and what each class is capable of imitating.

What this topic owns is the top of the scheme: the division of every seizure by where it begins, the reason that division is the first question asked, the closed list of causes that sits beneath it, and the single idea that earns all of this a place in a psychiatry chapter, that seizures and psychiatric episodes can impersonate one another. The finer subdivision of focal seizures by awareness, the anatomy and semiology of temporal and frontal lobe epilepsy, the reading of the electroencephalogram, the neurobiology that links limbic circuitry to behaviour, and the epidemiology of the epilepsy-psychiatry association are each developed in their own right elsewhere in this chapter; here they are named, not reopened.

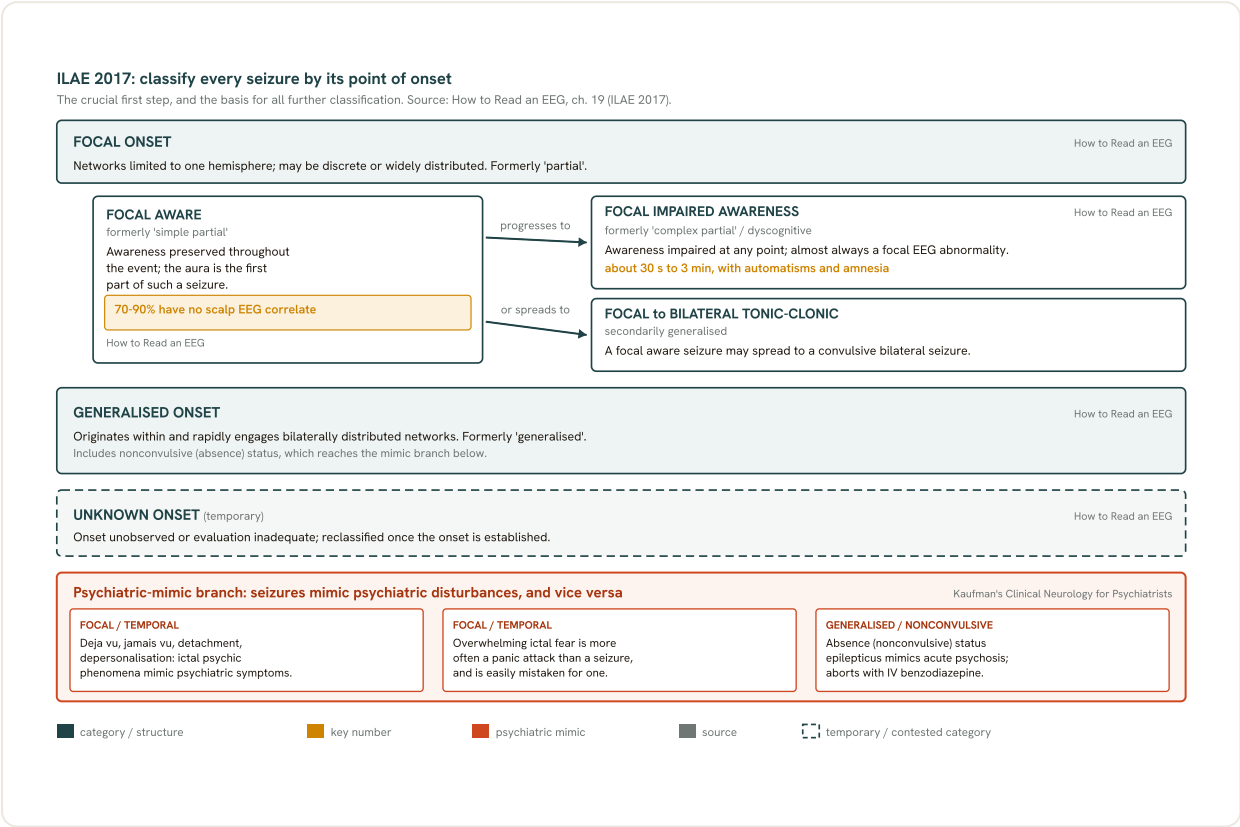


Fig 2.2 ILAE 2017 seizure classification. Every seizure is first sorted by point of onset into focal, generalised or unknown; focal onset is further split by whether awareness is preserved, with a focal aware seizure able to progress to impaired awareness or spread to a bilateral tonic-clonic seizure. The coral band collects the presentations that mislead the psychiatrist, since ictal phenomena mimic psychiatric disturbances and the reverse also holds.

2.1 The framework: seizures named by where they begin

A stepwise instrument, not a single label. The scheme works in tiers. The first tier concerns onset alone; only afterward does it ask what kind of epilepsy the person has, whether the picture fits a recognised epilepsy syndrome, and what is driving it. Naming the onset is therefore the doorway to a longer chain of reasoning, and it is not read off any one source. It is assembled from three streams held together: the semiology of the attack as it is witnessed or recounted, the electroencephalographic signature, and neuroimaging. None of the three is sovereign. A vivid eyewitness account with no electrical correlate, or a florid tracing with no history, leaves the onset a matter of inference rather than fact, and the classification is candid that its categories are provisional and sharpen as evidence accumulates. This is a genuine advance over older habits, which reached for a syndrome label first and let the description follow; anchoring the whole

edifice to the plain question of onset forces the clinician to earn each later refinement. The reading of the discharges themselves, ictal against interictal, is a craft of its own and is developed separately, as are the individual localisations that give focal epilepsies their names. What matters at this level is the shape of the system: a small set of onset families sitting above a short, closed list of causes, with the individual patient slotted in as the data allow. Holding that shape in mind stops the clinician leaping to a syndrome or a cause before the humbler question, where did it begin, has been honestly answered.

2.2 Focal onset: seizures that begin on one side

Confinement at the start is the defining feature. **Focal onset** seizures, the events once grouped together as partial seizures, begin within networks confined at their outset to a single hemisphere. That confinement is what defines them, but it should not be read as "small". The originating network may be tightly discrete, a single sharply localised generator, or it may be widely distributed across that hemisphere while still remaining on one side. Two seizures can both be focal in onset and yet look utterly different at the bedside, one a fleeting subjective symptom, the other a florid motor or behavioural event, because the size and situation of that one-hemisphere network vary enormously. What unites them is provenance, not appearance. This is also why focal onset is the family that carries most of the localising information a clinician can act on: because the seizure starts somewhere in particular, its earliest features often betray where, and it is those early, still-focal features, rather than any later spread, that point to a lobe and, with it, to a characteristic set of behavioural and experiential accompaniments. The subdivision of focal seizures by whether awareness is retained or impaired, and the specific psychiatric mimics that go with each, belongs to its own discussion, as do the individual localisations, the temporal and frontal presentations that dominate the psychiatric literature. Here the point is simply that a focal seizure is defined by its one-sided beginning and by nothing about how dramatic it looks.

2.3 Generalized onset: seizures that begin on both sides at once

Bilateral from inception, not a focal seizure that has spread. **Generalized onset** seizures are the mirror image: from the very beginning they arise within and rapidly recruit networks distributed across both hemispheres, so there is no single side of origin to find. The event is bilateral at its inception. This distinction is easy to state and genuinely hard to make at the bedside, because a focal seizure that propagates widely and a seizure that was generalized from the outset can converge on the same dramatic convulsion, and the observer who arrives late sees only the shared endpoint. The generalized family spans a wide clinical range, from brief lapses of awareness, through myoclonic jerks, to the major convulsion, and that range matters to psychiatry for one reason above all others: some generalized seizures are nonconvulsive. They produce no jerking and no fall, only an altered state, and a nonconvulsive generalized seizure that persists can present as a change in behaviour and awareness with nothing motor to declare it a seizure at all. That single possibility, a seizure with no convulsion, is where the classification stops being an academic exercise for the psychiatrist and becomes an urgent differential diagnosis, as the closing section of this discussion sets out in full.

2.4 Unknown onset: an honest placeholder

The third family is an admission, not a third kind of seizure. **Unknown onset** is used when the beginning of the event was not observed, or when the available evaluation is simply inadequate to place it, and it is explicitly a temporary designation, meant to be revised the moment better information arrives and the onset can be established. It exists because the honest alternative, forcing every event into focal or generalized on insufficient evidence, is worse, manufacturing a confidence the data do not support. A seizure that occurs in sleep with no witness, or a first event described only at third hand, may be genuinely unclassifiable at first contact, and the scheme would rather say so than guess. Importantly, an event can still be described by some of its features while its beginning stays unsettled, so "unknown onset" is not the same as "unknown event"; it is a statement about one axis only. The category is thus a discipline as much as a label, committing the clinician to returning to the question rather than closing it. It should never harden into a permanent verdict, because the whole point of naming a seizure is lost if the name is a shrug that no one revisits.

ONSET FAMILY	NETWORK AT THE START	STATUS OF THE LABEL	BEARING ON THE PSYCHIATRIC DIFFERENTIAL
Focal onset	One hemisphere	Definitive once established	Predicts a focal, experiential repertoire that can imitate panic or dissociation
Generalized onset	Both hemispheres at once	Definitive once established	Includes nonconvulsive states that can imitate acute psychosis
Unknown onset	Not established	Provisional, to be revised	Withholds the differential until the onset is known

The three onset families set against the axes that matter to a psychiatrist; the label's provisional or definitive status is itself part of the diagnosis.

2.5 The first fork: focal against generalized

One split does the heaviest clinical work. Of the three onset families, the division that carries the load is the split between focal and generalized, and it is worth pausing on why. The two describe genuinely different machinery. A seizure confined at its start to networks in one hemisphere behaves, spreads, responds and recurs unlike one that engages both hemispheres from the outset, and almost every subsequent decision keys off which of the two is in front of you. That is why the classification treats it as the pivot rather than a detail. Getting it right narrows the field of possible syndromes, points the causal work-up in a direction, and steers treatment toward what suits that seizure type; getting it wrong sends the entire work-up down the wrong branch, where every later step only lends false confidence to the original error. The fork determines, in turn:

- which epilepsy syndromes are even in contention;
- which causes justify a costly or invasive work-up;

- how the electrical and imaging investigation is targeted;
- which broad treatment pathway is rational for the seizure type.

RULE **Settle focal versus generalized before anything else.** The **focal-versus-generalized distinction** is the first and most consequential division in the whole scheme, because syndrome, cause and the entire work-up descend from it. Every later refinement is a branch of this first choice, which is why an unclassified seizure is only ever a half-diagnosed one.

TRAP **A witnessed bilateral convulsion is not proof of generalized onset.** A focal seizure can spread until both hemispheres are engaged and the patient convulses; what an observer sees at the peak says little about where the event began. Candidates equate the dramatic generalized motor phase with generalized onset and miss the focal beginning, which is exactly the beginning that carries the localising and psychiatric information.

Name the onset first: focal, generalized or unknown is the fork from which the syndrome, the cause and the psychiatric differential all descend.

2.6 The aetiological axis: cause as a separate question

A closed list of causes, considered after the type is settled. Once the seizure has an onset and the epilepsy has a type, the classification turns to cause, and here too the answer is drawn from a defined list rather than an open field. The **aetiological axis** sorts the epilepsies into **six** categories, structural, genetic, infectious, immune, metabolic and unknown, and cause is examined only after the seizure and epilepsy type are established, not before. The categories are not mutually exclusive in spirit, since a single patient may carry a structural lesion that is itself genetic in origin, but the axis forces the clinician to ask the question deliberately and to hold a place for each kind of answer. The sobering counterweight to all this machinery is how often it returns nothing: in surveyed populations the cause is never identified in about **two-thirds** of people with epilepsy, a fraction that holds broadly across age groups. That figure is not a failure of the scheme so much as an honest measure of the field, and it is why "unknown" is a named category rather than an embarrassed silence. For the psychiatrist the axis carries a specific warning: the structural, infectious, immune and metabolic causes are precisely the ones that also produce psychiatric symptoms directly, so a seizure filed into one of those causal boxes is a seizure whose psychiatry may be organic rather than reactive, and worth chasing to its root.

MNEMONIC

Onset first, cause second

Read every seizure in that order: fix where it begins before arguing about why it happens. The onset families sit above the closed list of causes, and jumping to aetiology before onset is settled is how a focal, structurally driven seizure gets mislabelled as a generalized epilepsy with no lesion to find.

PITFALL

Do not let the two "unknowns" blur into one. Unknown onset is a statement about the seizure, that we could not see where it began. Unknown cause is a statement about the epilepsy, that we could not find why it happens. They are independent: a seizure can have a confidently classified onset and still, like the majority of people with epilepsy, no identifiable cause at all. Writing "unknown" once and quietly meaning both is a classification error that buries real information.

2.7 Why the psychiatrist needs the classification: the logic of mimicry

The interface is a two-way street. The reason a scheme built by neurologists earns space in a psychiatry chapter is a single, unnerving fact: **psychiatric mimicry** runs in both directions. Seizures can counterfeit psychiatric disturbance, and psychiatric events can counterfeit seizures, so a clinician who cannot classify the first cannot safely diagnose the second. Focal seizures in particular can deliver brief experiential symptoms, fear, an altered sense of familiarity, a wave of unreality, that resemble a panic attack or a dissociative episode, while paroxysmal behavioural events with no electrical basis can arrive dressed as convulsions. The finer subdivision of focal seizures by preserved or impaired awareness, and the specific mimics that follow, is taken up separately, as are the ictal fear and panic phenomena that most closely imitate a psychiatric attack. Getting the onset family right is simply the first filter, because each family predicts a different repertoire of subjective events and therefore a different set of look-alikes. This is a different question from how often epilepsy and psychiatric disorder occur together and raise the risk of one another; that bidirectional comorbidity, and its epidemiology, is developed in its own right elsewhere in this chapter. The concern here is phenomenological rather than statistical: not how commonly the two coexist, but how convincingly one can pass for the other in a single patient on one occasion, and why the remedy is disciplined classification rather than pattern-matching.

IN INDIA

In much of Indian practice the psychiatrist is the first physician a family consults about a "behavioural spell", and the event itself is often unwitnessed by anyone who can describe it, with routine electroencephalography and especially video-electroencephalography unevenly available outside larger centres. More events therefore arrive already provisional, their onset genuinely undetermined rather than merely unrecorded, and more of them are carried for a long time under a psychiatric or a faith-healing label before a seizure is even considered. The discipline of classifying by onset, and of holding the category open until better evidence comes, matters more, not less, exactly where that evidence is hardest to obtain.

2.8 Absence status epilepticus: the seizure that impersonates psychosis

One generalized state is a near-perfect forgery of a psychiatric emergency. The clearest single reason a psychiatrist must own this material sits at the far end of the generalized family. **Absence status epilepticus**, the generalized nonconvulsive form, produces no dramatic convulsion at all. Instead the patient is altered: apathetic, confused, slowed, sometimes mute or oddly disinhibited, a state that can persist for **several hours** and reads convincingly as acute psychosis or an acute

confusional presentation. There is no aura to report and no fall to witness, so the history offers nothing to flag a seizure, and the referral arrives already labelled psychiatric. The distinguishing move is electrographic: the diagnosis is confirmed on the electroencephalogram, and the episode characteristically resolves once an intravenous benzodiazepine is given under that electrographic confirmation. That responsiveness is itself informative, because a primary psychosis does not lift within minutes of a benzodiazepine. The lesson generalises well beyond this one syndrome. A generalized nonconvulsive state can wear the face of psychiatry precisely because it strips away the motor signature that would otherwise announce a seizure, and it is only classification, knowing that a generalized onset can be nonconvulsive, that keeps the possibility alive in the clinician's mind. The peri-ictal psychiatric syndromes proper, sorted by their timing to the seizure, are taken up in their own right; here the point is narrower and prior, that the very existence of a nonconvulsive generalized seizure is a standing reason to classify before you diagnose.

GOOD TO REMEMBER

When a patient is sent to you for "acute psychosis" but the mental state is fluctuating, the affect blunted and the sensorium clouded, put nonconvulsive status on the list before you reach for an antipsychotic, and ask for an electroencephalogram. A single benzodiazepine given under electrographic cover can be both the test and the treatment, and getting there first changes the whole trajectory of the presentation.

2017	six	two-thirds
ILAE ONSET-BASED RECLASSIFICATION	AETIOLOGICAL CATEGORIES OF THE EPILEPSIES	EPILEPSY WITH NO IDENTIFIED CAUSE

3 · Focal aware versus focal impaired-awareness seizures and their psychiatric mimics

This is where the psychiatrist meets epilepsy first, and often without recognising it. A patient describes a recurring wave of unreality, a flood of familiarity in a strange room, a rising dread that peaks and passes in a minute, or brief lapses during which relatives say the person fumbled with their clothes and remembers nothing. Each of these is a psychiatric complaint on its face, and each can be the whole of a focal seizure. The examinable heart of this topic is the single distinction that organises focal seizures for the modern clinician, whether awareness is kept or lost, and the corollary that matters to us specifically: the symptoms produced can be indistinguishable at the bedside from panic, depersonalisation, dissociation and fugue. Reading the awareness axis correctly, and holding the psychiatric mimics in mind, is what separates a resident who commits a patient to years of the wrong treatment from one who asks the right next question.

Focal seizures are those that begin in one part of one hemisphere, and their overarching place within the broader division into focal, generalised and unknown-onset seizures is set out in its own topic and not reopened here. What this topic owns is the layer beneath that division: how a focal seizure is subclassified by awareness, what each subtype looks like and how long it runs, why the surface electroencephalogram behaves so differently across the two, how one grows into a

bilateral convulsion, and, throughout, the psychiatric conditions these events counterfeit. The anatomy and semiology of temporal lobe epilepsy, the substrate of most of these auras, are developed in their own right elsewhere in this chapter, as are frontal lobe epilepsy and its behavioural presentations; here they are named, not re-taught.

3.1 The single dividing line: awareness

One question splits the whole category. A focal seizure is classified by whether the patient is aware of themselves and their surroundings while it happens. A **focal aware seizure** is one in which awareness is preserved throughout the entire event, so the patient can later give an account of it from the inside; a **focal impaired-awareness seizure** is one in which awareness is disturbed at any point during the event, even briefly, even if it is regained before the seizure ends. Awareness here is used in a precise operational sense: the patient's connectedness to the event and their capacity to recall and respond during it, not merely whether they stay upright or their eyes stay open. The reason a single binary carries so much weight is that it maps onto everything the clinician then wants to know. The aware patient is a reliable narrator of their own aura and can warn of what is coming; the patient whose awareness is breached cannot, so the history has to be built from a witness, and the event carries a different risk to driving, to work and to safety. These two labels describe a state during a seizure, not two separate diseases: the same person can have both, and the same seizure can begin in one state and pass into the other.

RULE **Awareness, not motor drama, is the classifier.** A focal seizure is aware if consciousness of the self and surroundings is preserved from start to finish, and impaired-awareness if it is disturbed at any moment during the event. The presence or absence of convulsive movement, falling or eye-opening does not decide the class; the integrity of awareness does.

3.2 The names that were retired, and why the change matters

Old vocabulary, familiar to examiners, now superseded. The current terms replaced an older pair that still surfaces in wards, in question banks and in the memories of senior colleagues, so both must be recognised. What is now called a focal aware seizure was formerly the **simple partial** seizure, and what is now the focal impaired-awareness seizure was the **complex partial** seizure, a category some sources also labelled **dyscognitive**. The relabelling is not cosmetic. The words "simple" and "complex" described the seizure's apparent elaborateness, which is exactly the wrong axis, because a focal aware seizure can be a rich, frightening, deeply strange experience while a focal impaired-awareness seizure can look outwardly quiet. Anchoring the name to awareness rather than to how complicated the event seems forces the clinician to ask the one clinically decisive question and stops the reflexive assumption that a dramatic-looking spell must be the more serious category. For the psychiatrist the shift is doubly useful: the experiential auras that most concern us occur with awareness fully intact and sat awkwardly under a label like "simple".

TRAP **Reading "impaired awareness" as "unconscious".** Candidates equate the impaired-awareness seizure with loss of consciousness, coma or collapse. Impaired awareness need not mean unconscious: the patient may remain standing, eyes open, carrying out automatisms, and simply be disconnected from and later amnesic for the event. Describing it as a faint or a blackout misses the point and loses the mark.

3.3 Duration, automatisms and amnesia of the impaired-awareness seizure

A stereotyped shape: a short spell, purposeless activity, a hole in memory. The focal impaired-awareness seizure has a characteristic tempo and content that make it recognisable once the pattern is known. It typically lasts **about 30 seconds to 3 minutes**, a duration that is one of its most useful diagnostic anchors because it sits between the momentary and the prolonged. During the event the patient commonly shows purposeless automatisms with dense amnesia for what happened, and, depending on the region involved, transient disturbances such as aphasia may colour the picture. Automatisms are the involuntary, often repetitive fragments of behaviour that surface when higher control is offline: lip-smacking, chewing or swallowing, fumbling or picking at clothing, aimless walking, or repeating a snatch of speech. They look semi-purposeful precisely because they are assembled from over-learned motor programmes running without the awareness that would normally direct them, which is why the patient can appear to be doing something while being genuinely absent from it. The amnesia is not incidental; it is the behavioural signature of the awareness breach that defines the class. A witness who watched the whole spell is therefore worth more than the patient's own account, since by definition the patient did not fully register it. The regional flavour of the automatisms and language change is where the semiology of specific lobes begins, and the fuller anatomy belongs to the topics that own it; what matters here is the generic template of a brief spell, automatism and amnesia that flags the impaired-awareness class.

3.4 The electroencephalographic divergence

The scalp trace lies far more often for the aware seizure. One of the highest-yield and most counter-intuitive contrasts between the two subtypes is how they appear, or fail to appear, on the surface electroencephalogram. In the great majority of focal aware seizures there is no accompanying surface change at all: no scalp electroencephalographic correlate is found in **70 to 90 per cent** of them, so a completely normal routine trace recorded during a genuine focal aware seizure is the rule rather than the exception. The focal impaired-awareness seizure behaves in the opposite way: it almost always shows a focal ictal abnormality, a discharge or rhythmic ictal activity that can then spread to involve both hemispheres. The mechanism behind the difference is instructive and worth carrying. A focal aware seizure can arise from a discharge confined to a small volume of cortex, too limited in extent, or too deep, to project a detectable signal through the skull to scalp electrodes, whereas the loss of awareness that defines the other subtype requires recruitment of a larger, often bilateral network, and that broader synchronised activity is what the surface electrodes can finally see. The general principles of the electroencephalogram and of ictal versus interictal discharges are set out in their own topic; what this contrast establishes is a specific and clinically loaded fact about these two seizure types.

PITFALL

Concluding that a patient does not have epilepsy because a routine electroencephalogram recorded during, or between, their spells was normal. For focal aware seizures a normal scalp trace is expected, not reassuring, and even an ictal recording may show nothing. A normal electroencephalogram never excludes a focal seizure; the diagnosis rests on the clinical account, and the investigation is chosen to confirm, not to rule out.

3.5 Evolution: the aura and the march to a bilateral convulsion

The aware seizure is often only the first chapter. A focal aware seizure is not always a self-contained event. It may progress to impaired awareness, or spread further to a **focal-to-bilateral tonic-clonic seizure**, the event formerly called a secondarily generalised convulsion. This is the single most important reframing in the topic for a psychiatrist, because it dissolves the apparent gap between a strange subjective symptom and a full convulsion. The warning that patients and clinicians call an **aura** is not a herald that precedes the seizure; it is the first portion of the seizure itself, the focal aware phase experienced before awareness is lost or the discharge generalises. Understanding this changes how the history is heard. When a patient reports a consistent rising sensation, an odd smell or a wave of familiarity that on some occasions is followed by a convulsion and on others resolves on its own, they are not describing two problems but the same seizure captured at two lengths. It also explains why the aware phase is so precious diagnostically: it is the only part of the event the patient can narrate from within, and its content often points to where the seizure begins. The sequence is best held as a spectrum of extent rather than as separate diagnoses.

- The focal aware phase, the aura, in which the discharge is contained and awareness is intact.
- Progression to impaired awareness as the network recruited by the seizure widens.
- Spread to a focal-to-bilateral tonic-clonic seizure when the discharge generalises across both hemispheres.

3.6 Psychic and experiential auras as psychiatric mimics

Now the reason these seizures land in a psychiatry chapter at all. Because the focal aware phase is experienced and remembered, its content can be purely mental, and when it is, it counterfeits psychiatric illness with uncanny fidelity. Kaufman's Clinical Neurology for Psychiatrists describes the **experiential or psychic aura** as exactly this class of phenomenon: a subjective, often emotional or perceptual experience generated by the seizure. The mimics that recur are a small, memorable set. There is a sudden, unearned sense of familiarity, deja vu, and its mirror image, an alienating unfamiliarity within well-known surroundings, jamais vu. There is a feeling of unreality and self-estrangement that is clinically inseparable from **depersonalisation** and derealisation, and a sense of detachment from the self and the world. Presented in a clinic without the surrounding history, any of these reads as a dissociative symptom, a depersonalisation-derealisation complaint, or an anxiety phenomenon. Two features rescue the diagnosis. The first is form over content: an ictal experience tends to be brief, stereotyped and recurrent, and it may shade into a lapse or a convulsion in a way a primary psychiatric symptom

does not. The second is a discipline of suspicion, particularly where the same experience arrives in identical form again and again. One caution belongs with the enthusiasm, though: these experiential phenomena are not reliable for localisation, so their presence flags a possible seizure without pinning down where it starts.

MNEMONIC

BASE

What tilts a paroxysmal "mental" symptom away from primary psychiatry and toward a focal seizure. **B**rief, running seconds to a few minutes rather than a sustained state. **A**mnesia for part of the event once awareness is breached. **S**tereotyped, reproducing nearly the same experience each time. **E**pisodic and recurrent, arriving out of the blue. Each letter is a feature this topic states; none is decisive alone, and the base rates still matter.

3.7 Ictal fear and the panic-attack differential

The commonest and most consequential mimic of all. Of the experiential auras, the one that most often reaches a psychiatrist is a paroxysm of sudden, overwhelming fear, and it is the trickiest because it collides with one of our commonest diagnoses. A focal seizure, typically of temporal origin, whose entire content is a surge of **overwhelming fear** can be almost impossible to tell from a panic attack on phenomenology alone: both bring abrupt terror, autonomic arousal and a sense of impending doom, both peak fast and subside, and both leave the person shaken. The honest teaching point, and the one Kaufman's Clinical Neurology for Psychiatrists makes, is a probabilistic one: an isolated attack of overwhelming fear is more likely to represent a **panic attack** than a seizure, simply because panic is so much commoner, even as the two are genuinely easy to mistake for each other in both directions. The reflex should not be to relabel every panic patient as epileptic, but to keep ictal fear on the differential when the attacks are very brief, stereotyped, accompanied by any clouding of awareness or automatism, or followed on occasion by a convulsion. The ictal fear phenomenon itself, its amygdalar origin and the full range of ictal affective experience, is developed in its own topic and is not taught here; what this topic owns is the clinical differential with panic disorder, the psychiatric-mimic angle. Anxiety disorders in epilepsy, and the separation of an interictal anxiety disorder from an ictal one, are likewise handled in their own right elsewhere in this chapter.

3.8 The psychiatrist's discipline: reading a paroxysmal mental symptom

Put the axis and the mimics together at the bedside. The practical synthesis is to hold two questions at once for any patient with paroxysmal, unexplained mental phenomena. First, place the event on the awareness axis: was the person connected to and able to recall the episode throughout, which points to a focal aware event, or was there a breach of awareness with automatisms and amnesia, which points to the impaired-awareness class and mandates a witness account. Second, ask whether a symptom that presents as psychiatric, the depersonalisation, the deja vu, the wave of terror, might be the experiential content of a focal aware seizure rather than a primary disorder, and let the form of the episode guide the suspicion. Neither question is answered by a single investigation; the electroencephalogram in particular is expected to be silent

in the aware seizures that most often masquerade as psychiatric illness. The comparison that follows lays the two subtypes side by side on the axes that separate them, so the whole contrast can be weighed at once rather than feature by feature.

AXIS	FOCAL AWARE SEIZURE	FOCAL IMPAIRED-AWARENESS SEIZURE
Awareness	Preserved throughout the event	Impaired at some point during the event
Former term	"simple partial"	"complex partial", also dyscognitive
Consciousness	Fully retained	Impaired, though not necessarily lost
Automatisms and amnesia	Absent; the patient recalls the event	Purposeless automatisms with amnesia
Scalp EEG during the event	Usually no surface correlate	Almost always a focal ictal change
May evolve to	Impaired awareness or a bilateral convulsion	A focal-to-bilateral tonic-clonic seizure
Psychiatric mimic	Depersonalisation, dissociation, panic	Dissociative or fugue-like states

The two focal subtypes compared on the axes that separate them; awareness is the classifier, and the electroencephalographic and psychiatric-mimic rows are the ones that most often mislead.

30 s to 3 min	**70 to 90%**	**almost always**
FOCAL IMPAIRED-AWARENESS SEIZURE DURATION	FOCAL AWARE SEIZURES WITH NO SCALP EEG CORRELATE	FOCAL ICTAL CHANGE IN IMPAIRED-AWARENESS SEIZURES

IN INDIA

Two realities sharpen this topic in Indian practice. Access to routine and especially prolonged video-electroencephalography is concentrated in tertiary centres, so the aware seizures that carry no surface correlate are even more likely to be dismissed when the one available trace comes back normal; the diagnosis has to be carried on the clinical account. And where neurological services are distant, a patient whose only seizure is an experiential aura of deja vu, unreality or paroxysmal fear frequently enters care through a psychiatric or a faith-healing route and is labelled with an anxiety or dissociative disorder for a long time before a focal seizure is considered. A patient reporting stereotyped, brief, recurring mental symptoms deserves that question asked explicitly, exactly where confirmatory testing is scarce.

GOOD TO REMEMBER

When someone describes a recurring, identical, out-of-the-blue mental experience that lasts seconds to a couple of minutes, whether it is deja vu, a sense of unreality, or a surge of fear, add one question to the psychiatric history: does it ever end in a lapse of awareness, an automatism, or a convulsion, and does a witness describe absence or fumbling. A normal electroencephalogram does not close the matter. The awareness axis and the form of the episode, not the trace, carry the diagnosis.

Focal seizures are classified by awareness alone: preserved throughout in the focal aware seizure, breached at any point in the focal impaired-awareness seizure; the aware phase is the experienceable aura that mimics depersonalisation, deja vu and panic and that usually leaves the scalp electroencephalogram normal, so a brief, stereotyped, recurrent mental symptom is a possible seizure until the awareness axis and a witness say otherwise.

4 · Temporal lobe epilepsy: anatomy, semiology and psychiatric salience

No other epilepsy sits so squarely on the psychiatrist's desk. The temporal lobe is where seizure and mental phenomenon become the same event: an attack here can be a wave of fear, a flood of the uncanny, a phantom smell, a rising sickness in the gut, a minute of vacant fumbling the patient does not remember. Because these experiences are subjective and strange rather than convulsive, the person who has them is far more likely to reach a psychiatric clinic than a neurology ward, and far more likely to be told the trouble is anxiety, dissociation or a personality problem than to be recognised as having focal epilepsy. The marks here are for holding the anatomy, semiology and psychiatric consequences of **temporal lobe epilepsy** as one connected picture, so that an experiential aura is read as a possible seizure, a discrete and often operable lesion is looked for behind it, and the disproportionate psychiatric burden the temporal lobe carries is anticipated rather than discovered late.

The account moves from where these seizures come from, through what they look and feel like and how they mark the electroencephalogram, to why the temporal lobe is the seizure focus that psychiatry cannot ignore. The broader ILAE framework that names a seizure focal, generalised or of unknown onset, and the distinction between focal aware and focal impaired-awareness seizures, are developed in their own right elsewhere in this chapter. The interictal personality change this same epilepsy can produce, and the named Gastaut-Geschwind cluster, are constructs of their own and taught in their own right elsewhere in this chapter; this topic supplies the substrate on which they rest but does not teach them.

4.1 The temporal lobe as the home of impaired-awareness seizures

Where consciousness clouds, look first at the temporal lobe. The single anatomical fact from which everything else follows is that **focal impaired-awareness seizures** - the events an older nomenclature called complex partial seizures - arise most often from the temporal lobe, according to Kaufman's Clinical Neurology for Psychiatrists. The qualifier matters as much as the rule: the frontal lobe can also harbour the focus, and frontal lobe epilepsy with its own

behavioural and psychiatric presentations is a distinct topic taught in its own right elsewhere in this chapter. Within the temporal lobe the origin is not evenly spread. Kaufman's puts it precisely: **about 80%** of complex partial seizures originate from the medial, or mesial, surface of the temporal lobe rather than the outer convexity. That concentration is the reason electrodes placed close to the skull base, the nasopharyngeal and sphenoidal electrodes, were devised, since a discharge buried on the mesial surface may be too far from the standard scalp array to register clearly. The take-home is directional: a patient whose seizures involve altered awareness, automatism and amnesia has, until proven otherwise, a temporal, and probably mesial temporal, focus, and the diagnostic effort should be pointed there.

4.2 Mesial temporal sclerosis: the commonest substrate

The lesion behind the syndrome. If most temporal seizures begin on the mesial surface, the next question is what is wrong with that tissue, and here temporal lobe epilepsy has an answer unusually specific for a seizure disorder. Kaufman's Clinical Neurology for Psychiatrists names the commonest cause as **mesial temporal sclerosis**, a combination of sclerosis of the hippocampus with atrophy of the temporal lobe. This is a discrete, identifiable and often surgically remediable pathology, which is why establishing it is worth the trouble: unlike a diffuse encephalopathy, a sclerotic hippocampus can in selected patients be removed, and the seizures with it. The origin of the sclerosis carries an examinable subtlety that is easy to garble. Kaufman's is explicit that it is prolonged febrile seizures, not the brief occasional ones, that are implicated in mesial temporal sclerosis; the association is strong but the direction of causation is still debated, and the ordinary short febrile convulsion of childhood does not, as a rule, scar the hippocampus. The same source implicates perinatal anoxia and herpes simplex encephalitis among the other insults that leave this signature. A long, complicated febrile seizure in early childhood is thus a meaningful pointer toward later mesial temporal lobe epilepsy, whereas a run-of-the-mill febrile fit is not, and conflating the two is a classic error.

TRAP **Attributing mesial temporal sclerosis to any febrile seizure history.** The examiner offers a patient with temporal lobe epilepsy and a childhood of febrile convulsions and invites an indiscriminate link. The link is real but conditional: the prolonged, complicated febrile seizure is implicated in mesial temporal sclerosis, while brief, occasional febrile seizures are not. Counting every febrile fit as a cause loses the distinction the question tests; denying the association altogether loses the other half of it.

4.3 The mesial aura: the experiential front door

The seizure that arrives as a feeling. Mesial temporal seizures are mistaken for psychiatric events because their onset is subjective and non-motor. How to Read an EEG describes a mesial temporal seizure as beginning with a nonmotor onset and behavioural arrest, the patient falling still and inaccessible, preceded or accompanied by an aura drawn from a small and recognisable family. There is the **rising epigastric sensation** and its autonomic accompaniments, an ascending visceral feeling that patients struggle to name. There is the **olfactory or gustatory aura**, a hallucinated smell or taste, usually unpleasant. There is the cognitive or experiential aura, the intrusions of **déjà vu and jamais vu**, the false familiarity of the new or the false strangeness

of the known. And there is the emotional aura, a sudden, motiveless fear welling up without a frightening thought behind it. Each is routinely misread as a primary psychiatric symptom: the epigastric rush as a panic attack, the fear as an anxiety disorder, the déjà vu and depersonalised strangeness as dissociation, the phantom smell as a psychotic hallucination. The dedicated treatment of ictal fear and ictal panic as amygdalar phenomena, and of the psychiatric mimics of focal seizures, belongs to topics of their own elsewhere in this chapter; the fact to hold is that these auras are the semiology of the mesial temporal lobe, and their psychiatric costume is what makes them dangerous.

MNEMONIC

FROG

The mesial temporal aura families in four beats, each one a feature How to Read an EEG actually names. **F**ear, rising without a reason. **R**ising epigastric sensation with its autonomic tail. **O**lfactory or gustatory hallucination, a smell or taste that is not there. **G**nostic disturbance, the déjà vu and jamais vu of a mind briefly unsure what it knows.

4.4 Automatisms and the vacant minute

The motor signature is quiet, repetitive and unremembered. Once awareness is impaired, the temporal seizure fills the gap with automatism rather than the violent movement of a frontal attack. Kaufman's Clinical Neurology for Psychiatrists fixes the frequency: **automatisms** are present in **about 80%** of seizures originating in the temporal lobe, making them the rule rather than the exception. They fall into two families the source names directly. The oral automatisms are the lip-smacking, chewing and swallowing movements of the mouth; the manual automatisms are the fumbling, picking and fidgeting of the hands. To a bystander this looks like a person who has drifted off and is aimlessly busy, not like a seizure, and afterwards the patient has no memory of the interval. The mimicry deepens when a minority of patients utter brief, stereotyped phrases, so that a fragment of speech emerges from someone who is not, in any usual sense, present. The yield is that a witnessed spell of altered contact filled with oro-alimentary and hand automatisms, followed by amnesia, is temporal lobe epilepsy until proven otherwise, however psychological the surface behaviour appears.

4.5 The electroencephalographic signature of the mesial focus

The record points to the anterior temporal region. The EEG in mesial temporal lobe epilepsy has a recognisable grammar that converts a suggestive history into a localised diagnosis. How to Read an EEG describes the interictal recording as showing anterior temporal discharges maximal at the F7 and F8 electrodes, and, in a proportion of patients, **temporal intermittent rhythmic delta activity**, or TIRDA, a run of rhythmic slow waves over the temporal region that is a useful lateralising sign. When a seizure is captured, the same source describes an ictal pattern of rhythmic alpha-theta activity or spikes over the involved temporal region, and a postictal record that commonly shows slowing on the same side as the focus, a further lateralising clue in the aftermath. The general principles of the EEG, and the difference between ictal and interictal discharges as they inform psychiatric assessment, are set out in their own right elsewhere in this

chapter; what belongs here is the fingerprint of the mesial temporal focus. That fingerprint is intermittent, which is the source of a common and costly misreading.

PITFALL

A single normal routine interictal EEG is taken to exclude temporal lobe epilepsy. The anterior temporal discharges are intermittent, and temporal intermittent rhythmic delta activity may only declare itself in drowsiness or sleep, so a brief waking recording can fall between the discharges and read as clean. Treating that clean trace as a rule-out abandons a patient with genuine, and potentially operable, mesial temporal epilepsy. The discipline is to repeat the study, capture drowsiness and sleep, and weigh the record against the semiology rather than let one normal trace close the question.

4.6 Lateral neocortical temporal epilepsy: the other half of the lobe

Not all temporal seizures are mesial. The examinable counterpoint is the lateral, or neocortical, temporal focus, which differs in aura, electrode localisation and behaviour, and holding the two apart is often the whole point of a question. How to Read an EEG characterises **lateral neocortical temporal lobe epilepsy** by an aura that is auditory or vertiginous rather than visceral or experiential: a hallucinated sound or a sense of spinning, reflecting the auditory cortex and its neighbours on the lateral surface. The interictal discharges shift with the anatomy, appearing over the mid and posterior temporal electrodes, T7 and T8, rather than the anterior temporal maxima of the mesial form. And the seizures spread differently: the same source notes a tendency to more frequent and more rapid secondary generalisation, so a lateral temporal seizure is more apt to march quickly into a bilateral convulsion than its mesial counterpart, which stays local longer. An auditory or vertiginous aura should therefore move the clinician's mental map from the mesial structures to the lateral neocortex and raise the expectation of quicker generalisation, changing both what is looked for on the EEG and what is warned about clinically.

FEATURE	MESIAL (MEDIAL) TEMPORAL	LATERAL (NEOCORTICAL) TEMPORAL
Characteristic aura	Rising epigastric, olfactory or gustatory, déjà vu or jamais vu, fear	Auditory or vertiginous
Onset character	Nonmotor onset with behavioural arrest	Sensory aura referable to the lateral surface
Interictal discharge maximum	Anterior temporal, F7 and F8; TIRDA	Mid and posterior temporal, T7 and T8
Spread to a bilateral convulsion	Tends to remain local longer	More frequent and more rapid secondary generalisation

The two temporal syndromes on their shared axes: the mesial focus announces itself through visceral and experiential auras and anterior temporal discharges, the lateral neocortical focus through auditory or vertiginous auras, more posterior discharges, and a readier march to generalisation.

4.7 Why the temporal lobe is psychiatry's seizure focus

The disproportionate mental burden. The temporal lobe earns its central place because, more than any other seizure origin, it is bound up with psychiatric illness, and the tightest of those bonds is with psychosis. The schizophrenia-like psychosis of epilepsy - the chronic interictal psychosis that can, over years, come to resemble schizophrenia - is especially strongly associated with temporal lobe epilepsy, a link brought into focus by the work of Slater and reported in Lewis Child and Adolescent Psychiatry. The entity itself, its full clinical features, its relation to the epilepsy and its management, is a construct owned by its own topic and taught in its own right elsewhere in this chapter; what this topic asserts is the salience, that of all the localisations it is the temporal lobe whose seizures most consistently keep company with this psychosis. That salience ties the earlier anatomy to a psychiatric outcome: the same mesial structures whose sclerosis generates the seizures sit within the limbic system that governs affect, memory and meaning, so their chronic irritation may express itself, in some patients, as disordered belief and perception. The functional link between the limbic system, kindling and behaviour is the mechanism step of this chapter and is taught in its own right elsewhere in it; the point to carry is that a temporal seizure focus is a psychiatric risk factor, and a patient with longstanding temporal lobe epilepsy who develops a psychosis is showing a recognised association, not a coincidence.

The same lobe is also the substrate of the interictal personality change and the eponymous behavioural cluster that longstanding temporal epilepsy can bring, both taught as constructs of their own elsewhere in this chapter and not reopened here. Together they make the temporal lobe the one seizure focus a psychiatrist must know in anatomical detail: its seizures wear psychiatric clothes at onset and its chronic course carries psychiatric consequences at the far end.

GOOD TO REMEMBER

When a patient presents with unexplained episodic fear, a rising sensation in the stomach, recurrent déjà vu or a phantom smell, do not close on a primary anxiety, panic or dissociative label without asking whether these are auras; a witnessed spell of altered awareness with lip-smacking or fumbling, followed by amnesia, points the same way. The move that changes the course is to think temporal lobe epilepsy, take a history for prolonged childhood febrile seizures, and arrange neurological assessment with an EEG that captures drowsiness.

RULE **An experiential or visceral aura is a temporal lobe seizure until proven otherwise, and its commonest substrate is a discrete lesion.** Most focal impaired-awareness seizures are temporal in origin, most temporal seizures begin on the mesial surface, and the commonest thing wrong with that surface is mesial temporal sclerosis. The chain matters because its end point is remediable: a subjective, easily psychologised aura can trace back to a sclerotic hippocampus that surgery may cure, which is why the aura must be taken seriously as neurology and not filed as psychiatry.

IN INDIA

Mesial temporal sclerosis is among the most surgically remediable of the epilepsies, yet the investigations that identify it, dedicated epilepsy-protocol MRI and prolonged or video-EEG, are concentrated in tertiary and academic centres, while most patients are first seen where only a routine interictal recording, or none, is available. A patient whose experiential auras and automatisms are the face of an operable lesion is thus often carried for years as treatment-resistant anxiety, dissociation or psychosis. The corrective is a low threshold to elicit the aura and the childhood febrile-seizure history, to name temporal lobe epilepsy, and to refer toward the imaging and monitoring that can turn a chronic psychiatric label into a treatable neurological diagnosis.

4.8 Recognising temporal lobe epilepsy: a practical synthesis

The diagnosis is assembled from convergent clues, not clinched by one test. No single feature is decisive, and the surface EEG can miss the mesial focus between attacks; what secures the recognition is the way the strands run together and against the psychological first impression they create. Each of the following should move temporal lobe epilepsy up the differential rather than let it slide off.

- An onset that is a feeling: a rising epigastric sensation, a phantom smell or taste, déjà vu or jamais vu, or motiveless fear, rather than a movement.
- A vacant interval of impaired awareness filled with oral and manual automatisms and followed by amnesia for the episode.
- A childhood history of a prolonged, complicated febrile seizure, which points toward mesial temporal sclerosis where a brief febrile fit would not.
- An interictal EEG showing anterior temporal discharges or temporal intermittent rhythmic delta activity, remembering that a single clean waking trace does not exclude the diagnosis.
- An auditory or vertiginous aura with more posterior temporal discharges when the focus is lateral rather than mesial.
- A longstanding course in which a psychosis strongly associated with temporal lobe epilepsy may emerge, marking the same lobe as the psychiatric one.

Read together, these turn a presentation that reads as anxiety, dissociation or a personality problem into one that reads as focal epilepsy with a nameable, often operable substrate. The clinician who has internalised the pattern treats an aura as neurology to be investigated rather than psychology to be soothed, a normal routine EEG as a reason to record again in sleep rather than to close the case, and a temporal seizure focus as a standing psychiatric risk to be watched, not a finished neurological fact to be filed away.

~80% OF COMPLEX PARTIAL SEIZURES ARE MESIAL TEMPORAL IN ORIGIN	**~80%** OF TEMPORAL-LOBE SEIZURES SHOW AUTOMATISMS	**mesial temporal sclerosis** COMMONEST SUBSTRATE OF TLE	**F7 / F8** ANTERIOR TEMPORAL INTERICTAL MAXIMUM

Temporal lobe epilepsy is the commonest source of focal impaired-awareness seizures, most of them mesial in origin and most often built on mesial temporal sclerosis; it announces itself through experiential and visceral auras and oro-alimentary automatisms rather than convulsion, and it carries the strongest link of any seizure focus to the schizophrenia-like psychosis of epilepsy, so its psychologically disguised onset must be read as neurology.

5 · Frontal lobe epilepsy and its behavioural and psychiatric presentations

This is the great impostor among the epilepsies. More often than any other seizure disorder, **frontal lobe epilepsy** reaches the psychiatrist before it reaches the neurologist, because its attacks look like almost anything except epilepsy. They erupt at night, throw the body into violent thrashing or rigid posturing, last only seconds, leave the patient clear-headed almost at once, and are accompanied by an electroencephalogram that is frequently, maddeningly, normal. To an eye trained on the slow, aura-led, confusion-trailed temporal lobe seizure, none of this reads as a fit. It reads as a nightmare, a tantrum, a functional display, or an act. The marks in this topic are awarded for resisting exactly that first impression, because the price of yielding to it is a patient with real, treatable epilepsy who is instead labelled hysterical, psychogenic or a fabricator and left unprotected.

The reason the mistake is so easy is that every feature which should signal a seizure is either absent or actively misleading. The behaviour is bizarre rather than orderly; the duration is too short to resemble a classical convulsion; awareness is often preserved, which the untrained observer takes as proof the patient is acting deliberately; and the surface EEG, the one investigation that ought to settle the matter, commonly shows nothing or is buried under the noise of the movements themselves. This account sets out the semiology, the timing, the electrographic pitfalls and the behavioural presentations of frontal lobe seizures so that the pattern, once seen whole, becomes recognisable rather than deceptive. The framework that separates focal aware from focal impaired-awareness seizures, and the fundamentals of the EEG, are developed in their own right elsewhere in this chapter; here the concern is what the frontal lobe in particular does, and why it fools people.

5.1 The motor semiology: hyperkinetic thrashing and tonic posturing

Dramatic movement is the signature. Frontal lobe seizures announce themselves through their motor content, and it is the character of that movement, not any accompanying inner experience, that defines them. How to Read an EEG describes two dominant patterns. In the first, the patient is thrown into **hyperkinetic** activity: agitated, complex and often forceful movements of the limbs and trunk, such as thrashing, kicking, bicycling, pelvic thrusting and rocking, which can look purposeful and are readily mistaken for a rage outburst or an emotional scene. In the second, one or more limbs are locked into **asymmetric tonic posturing**, a sustained and often bizarre stiffening in which the body is held in an unnatural attitude. Neither pattern resembles the slowly evolving oro-alimentary automatisms of the temporal lobe, and both look far stranger than a clinician expects a seizure to look. Just as important, awareness during these attacks may be preserved or lost, and it is the combination of odd, vigorous, apparently coordinated movement

with a patient who can sometimes still respond that most powerfully suggests the behaviour is willed. The formal distinction between seizures with and without preserved awareness belongs to the focal-seizure classification set out elsewhere in this chapter; what matters here is that retained awareness does not argue against a frontal seizure, and must never be pressed into service as an argument for a psychogenic one.

5.2 The timing: nocturnal, out of non-REM sleep, and clustering

They are creatures of the night. A second defining feature is when the seizures happen. How to Read an EEG emphasises that frontal lobe seizures are characteristically **nocturnal**, arising specifically out of non-REM sleep rather than from wakefulness, and that they tend to occur in clusters rather than as isolated events. This fact does enormous diagnostic damage in the wrong direction. A patient who, night after night, sits up, thrashes, vocalises, postures and then settles, only to do it again, presents to the world as someone with a sleep disorder or a night-time behavioural disturbance, not as someone with epilepsy. The nocturnal, sleep-bound pattern is precisely what aligns frontal seizures with the parasomnias in the mind of an observer, and the clustering deepens the impression of a recurring nocturnal event rather than a scattered seizure tendency. For the psychiatrist the practical yield is high: any history of repeated, brief, stereotyped nocturnal episodes emerging directly from sleep should raise frontal lobe epilepsy explicitly, because the sleep-relatedness that seems to argue against a seizure is in fact one of its strongest fingerprints.

5.3 The temporal profile: brief, aura-free, and quick to clear

On and off, with nothing before and little after. Kaufman's Clinical Neurology for Psychiatrists lays out the course that separates a frontal seizure from a temporal one at the bedside. Frontal lobe seizures are brief, typically lasting **under one minute**; they are not heralded by a premonitory aura; and they are followed by little or no postictal confusion, so the patient returns to normal almost as soon as the movement stops. Each of these is the mirror image of the temporal lobe seizure, whose anatomy, semiology and psychiatric salience are taught in their own right elsewhere in this chapter: the temporal seizure classically opens with an experiential or visceral aura, unfolds more slowly, and trails a period of postictal clouding. The abruptness at both ends is one of the features most often misread: an event that starts without warning and stops as suddenly, leaving a lucid patient, does not fit the naive template of a seizure and is easily taken for a behavioural outburst the patient could have contained. Kaufman's also fixes the epidemiological frame that completes the picture: frontal lobe seizures tend to begin in adult life and to recur relatively often, on the order of **several times a month**. An adult with frequent, brief, aura-free nocturnal spells is the prototypical presentation, and it walks into psychiatric and sleep clinics far more readily than into a neurology ward.

5.4 The electroencephalographic problem: a seizure the EEG hides

The one test you trust is the one that fails you. The feature that turns frontal lobe epilepsy from a difficult diagnosis into a treacherous one is its behaviour on the EEG. Kaufman's Clinical Neurology for Psychiatrists is explicit that frontal seizures generate paroxysmal discharges which are hard to capture: the interictal recording is often entirely normal, and even the ictal recording,

taken while the seizure is actually in progress, is frequently **obscured by movement artifact**, since the very thrashing that characterises the attack drowns the electrical signal in muscle and motion noise. The anatomy compounds the problem: much of the frontal cortex lies on mesial and basal surfaces, folded away from scalp electrodes, so discharges there may never reach the surface at readable amplitude. A patient whose seizure is genuine but whose EEG is silent is the single most dangerous configuration here, because a normal EEG is precisely what an observer reaches for to declare the episodes non-epileptic. Kaufman's does offer a way through: repeated recordings may eventually reveal bifrontal or midline discharges, so persistence with the EEG, and recourse to prolonged or sleep-capturing techniques, is often what finally secures the diagnosis. The lesson is not that the EEG is useless but that a single normal study proves nothing, and the burden of proof must not be laid on an investigation this particular epilepsy is built to evade.

RULE **A normal interictal EEG does not exclude frontal lobe epilepsy.** Of all the localisation-related epilepsies, this is the one in which the routine scalp recording is most likely to be clean between attacks and swamped by artifact during them. A negative EEG is therefore worthless as a rule-out here, and the diagnosis must rest on the clinical pattern, with prolonged, sleep-capturing or video recording pursued rather than the single normal trace accepted.

5.5 False localisation and false generalisation

Even when the EEG speaks, it may lie about where. When a frontal discharge does break through to the surface, a further trap opens, and it is one of the most examinable points in the topic. How to Read an EEG warns that a frontal ictal pattern may appear to arise from the temporal lobe on the same side as the true focus, a false localisation to the ipsilateral temporal region; or it may spread so quickly across the midline that the record looks generalised from the outset, a phenomenon of **secondary bilateral synchronisation**. Both errors point the interpreter away from the frontal lobe where the seizure actually begins. The clinical stakes are considerable, because localisation drives the choice of imaging, the interpretation of the surgical work-up, and ultimately the decision about which region, if any, might be resected; a patient whose frontal seizures are read as temporal or as primarily generalised is investigated and counselled along an entirely wrong axis. The defence is the same discipline the whole topic demands: the electrographic impression must be weighed against the semiology and the timing, never taken at face value, because in the frontal lobe the surface record is prone to misdirect both about whether a seizure is present and about where it is coming from.

PITFALL

A scalp EEG that captures a frontal seizure but appears to localise it to the ipsilateral temporal lobe, or to be generalised from onset, can steer the entire work-up wrong. The interpreter, trusting the trace, sites the focus in the wrong place, and the imaging, the presurgical evaluation and the counselling follow the error. When the semiology is unmistakably frontal, a temporal or generalised-looking ictal pattern should prompt suspicion of false localisation or secondary bilateral synchronisation rather than a revised anatomical diagnosis.

5.6 Bizarre behaviour and the question of aggression

The most behaviourally florid of the focal epilepsies. The frontal lobe earns its place in a psychiatry chapter because of what its seizures do to behaviour. Kaufman's Clinical Neurology for Psychiatrists makes the sharp point that frontal lobe seizures, even set beside temporal lobe seizures, are the most apt to produce genuinely **bizarre behaviour** and the rare instances of aggressive violence seen in the ictal state. This is not a licence to read violence into every frontal seizure; ictal aggression remains uncommon, and when it occurs it is crude, brief and undirected rather than organised. The general subject of aggression, irritability and episodic dyscontrol in epilepsy, and the vexed forensic question of criminal responsibility, is a construct in its own right and is taught elsewhere in this chapter. The point specific to the frontal lobe is directional and comparative: of all the seizure localisations, this is the one from which strange, disinhibited and occasionally aggressive ictal behaviour is most likely to spring, more so even than the temporal lobe with which it is usually contrasted. That comparative primacy is why frontal seizures so often land in psychiatric hands: the behaviour is loud, the movement is odd, and the whole presentation begs to be read psychologically, when the discipline of the topic is to remember that a seizure can produce all of it.

MNEMONIC

FRONTAL

The signature of the frontal lobe seizure in seven beats, each one a feature the topic actually states. **F**requent, several times a month. **R**apid and brief, over in under a minute. **O**ut of non-REM sleep, so characteristically nocturnal. **N**o aura to warn of onset. **T**onic asymmetric posturing, or hyperkinetic thrashing, as the motor content. **A**wareness that may be kept or lost. **L**ittle or no postictal confusion, clearing at once.

5.7 The great mimic: psychogenic seizures and sleep disorders

Where all the threads converge. Everything set out so far, the bizarre movement, the preserved awareness, the abrupt on and off, the nocturnal timing and the silent EEG, gathers into a single clinical hazard. Kaufman's Clinical Neurology for Psychiatrists states it directly: because their behavioural manifestations are so bizarre and their EEG changes usually undetectable, frontal lobe seizures are the epileptic events most readily mistaken for **psychogenic non-epileptic seizures**, and, when they arise at night, for a sleep disorder. This is the psychiatric heart of the topic. Psychogenic non-epileptic seizures, their semiology, their investigation with video-EEG and their management, are defined and taught across their own dedicated topics elsewhere in this chapter; the fact that belongs here is the reverse direction, that frontal lobe epilepsy is the organic diagnosis most often lost to a functional label. The asymmetry of harm is what should be held: to call a psychogenic attack epileptic risks needless drugs, whereas to call a frontal seizure psychogenic risks abandoning treatable epilepsy and, with it, the seizure control on which so much else depends. A confident functional or sleep-disorder label placed on a bizarre nocturnal episode, on the strength of a normal routine EEG, is the archetypal error.

TRAP **Calling a frontal seizure psychogenic because it looks bizarre and the EEG is clean.** The examiner's favourite version of this topic offers a young adult with brief, dramatic, hyperkinetic or posturing nocturnal spells, preserved awareness, rapid recovery and a normal routine EEG, and invites the candidate to diagnose a non-epileptic or functional attack. Every one of those features is a positive fingerprint of frontal lobe epilepsy. The bizarre look and the normal EEG are not evidence against a seizure; in the frontal lobe they are evidence for one.

FRONTAL SEIZURE FEATURE	WHAT IS SEEN	WHY IT MISLEADS
Onset and frequency	Adult onset, recurring several times a month	An established, repetitive pattern reads as a behavioural or sleep habit
Duration and course	Under a minute, no aura, clears at once	Abrupt on and off looks non-physiological and suggests a willed act
Timing	Nocturnal, out of non-REM sleep, clustering	Sleep-bound recurrence mimics a parasomnia
Semiology	Hyperkinetic thrashing or asymmetric posturing	Odd, forceful movement is read as functional or emotional
Awareness	May be preserved through the attack	Retained awareness with strange movement suggests volition
Interictal EEG	Frequently normal	A clean record is wrongly taken to exclude epilepsy
Ictal EEG	May falsely localise or appear generalised	Even the ictal trace can point to the wrong lobe

The diagnostic signature of the frontal lobe seizure: each feature that ought to confirm a seizure is instead the feature that pushes the unwary observer toward a psychogenic or sleep-disorder diagnosis.

IN INDIA

The equipment that resolves this diagnosis is unevenly distributed. Prolonged and video-EEG monitoring is concentrated in tertiary and academic centres, while most patients are first seen where only a routine interictal recording is available. A single normal trace from a district setting cannot exclude frontal lobe epilepsy, yet it is often read as if it can, and the bizarre nocturnal spells that follow are relabelled as a psychogenic disturbance, a sleep problem or a primary psychiatric complaint. Sleep-related and behaviourally florid nocturnal seizures should be actively considered rather than assumed functional, and referral for prolonged or video-EEG, rather than reassurance drawn from a clean routine study, is what protects a patient with genuine, treatable epilepsy from being written off.

5.8 Recognising the pattern: a practical synthesis

The diagnosis is made by pattern, not by any single test. No single feature is decisive on its own, and the surface EEG, the investigation a clinician instinctively leans on, is the least reliable arbiter in this setting. What secures the diagnosis is the convergence of the features, read together and against the grain of the first impression they create. Each of the following should move frontal lobe epilepsy up the differential rather than down it.

- Repeated, brief, stereotyped nocturnal episodes that emerge directly out of sleep and tend to cluster.
- Hyperkinetic thrashing or fixed asymmetric posturing as the motor content, however bizarre it appears.
- Attacks that begin without an aura and clear at once, without the trailing confusion of a temporal seizure.
- Preserved awareness during obviously abnormal movement, which argues neither way and must not be read as proof of volition.
- A normal or unrevealing routine EEG that has been offered, wrongly, as reassurance.

Held together, these turn a presentation that reads as psychiatric or functional into one that reads as neurological, by refusing to let any single misleading feature carry the decision. The clinician who has internalised the pattern treats a normal EEG as a reason to record for longer rather than to close the case, treats preserved awareness as neutral rather than as proof of pretence, and treats the very strangeness of the behaviour as a pointer toward the frontal lobe rather than away from epilepsy. That reorientation is what keeps a treatable epilepsy from being mislaid among the disorders it so faithfully imitates.

GOOD TO REMEMBER

When a patient gives a history of repeated, brief, bizarre nocturnal spells and arrives with a normal routine EEG, do not let the clean trace close the door. Keep frontal lobe epilepsy on the differential, ask specifically about sleep-bound clustering, preserved awareness and rapid recovery, and arrange prolonged, sleep-capturing or video-EEG through neurology rather than settling for a functional or sleep-disorder label. The normal EEG is the trap, not the answer.

under 1 min	NREM sleep	several / month	often normal
TYPICAL SEIZURE DURATION	WHEN THEY ARISE	USUAL FREQUENCY IN ADULTS	INTERICTAL SCALP EEG

Frontal lobe seizures are brief, aura-free, nocturnal attacks of hyperkinetic thrashing or asymmetric posturing that clear at once and sit on a frequently normal interictal EEG; that pairing of bizarre behaviour with an unrevealing record is why they are the epileptic events most often misread as psychogenic or as a sleep disorder, and why a normal EEG must never be allowed to exclude them.

6 · Epidemiology and bidirectional relationship of epilepsy and psychiatric disorders

The interface begins with counting. Before the peri-ictal syndromes, the psychoses and the prescribing dilemmas of the topics that follow, the psychiatrist needs a working sense of how common epilepsy is, how it behaves across a life, and how tightly it is bound to psychiatric illness. Two ideas carry most of the marks in this area, and both reward being held clearly. The first is that epilepsy is common, chronic and, for the majority, remitting, so that its frequency at a single

moment and its frequency across a lifetime are genuinely different quantities answering different questions. The second, and the one candidates most often underrate, is that the link between epilepsy and psychiatric disorder runs in both directions: seizures raise the risk of psychiatric illness, and psychiatric illness raises the risk of seizures. Together these turn a short list of prevalence figures into a way of thinking about the shared ground where neurology and psychiatry meet.

The examinable content divides into the numbers a candidate is expected to quote and place in the right conceptual slot, and the two-way relationship whose logic reorganises how a clinician reads the patient in front of them. Drawing on Kaufman's Clinical Neurology for Psychiatrists, Lewis's Child and Adolescent Psychiatry and the Maudsley Prescribing Guidelines, it is the chapter's hinge: the peri-ictal, psychotic, mood and drug topics that follow all assume the frame set out here.

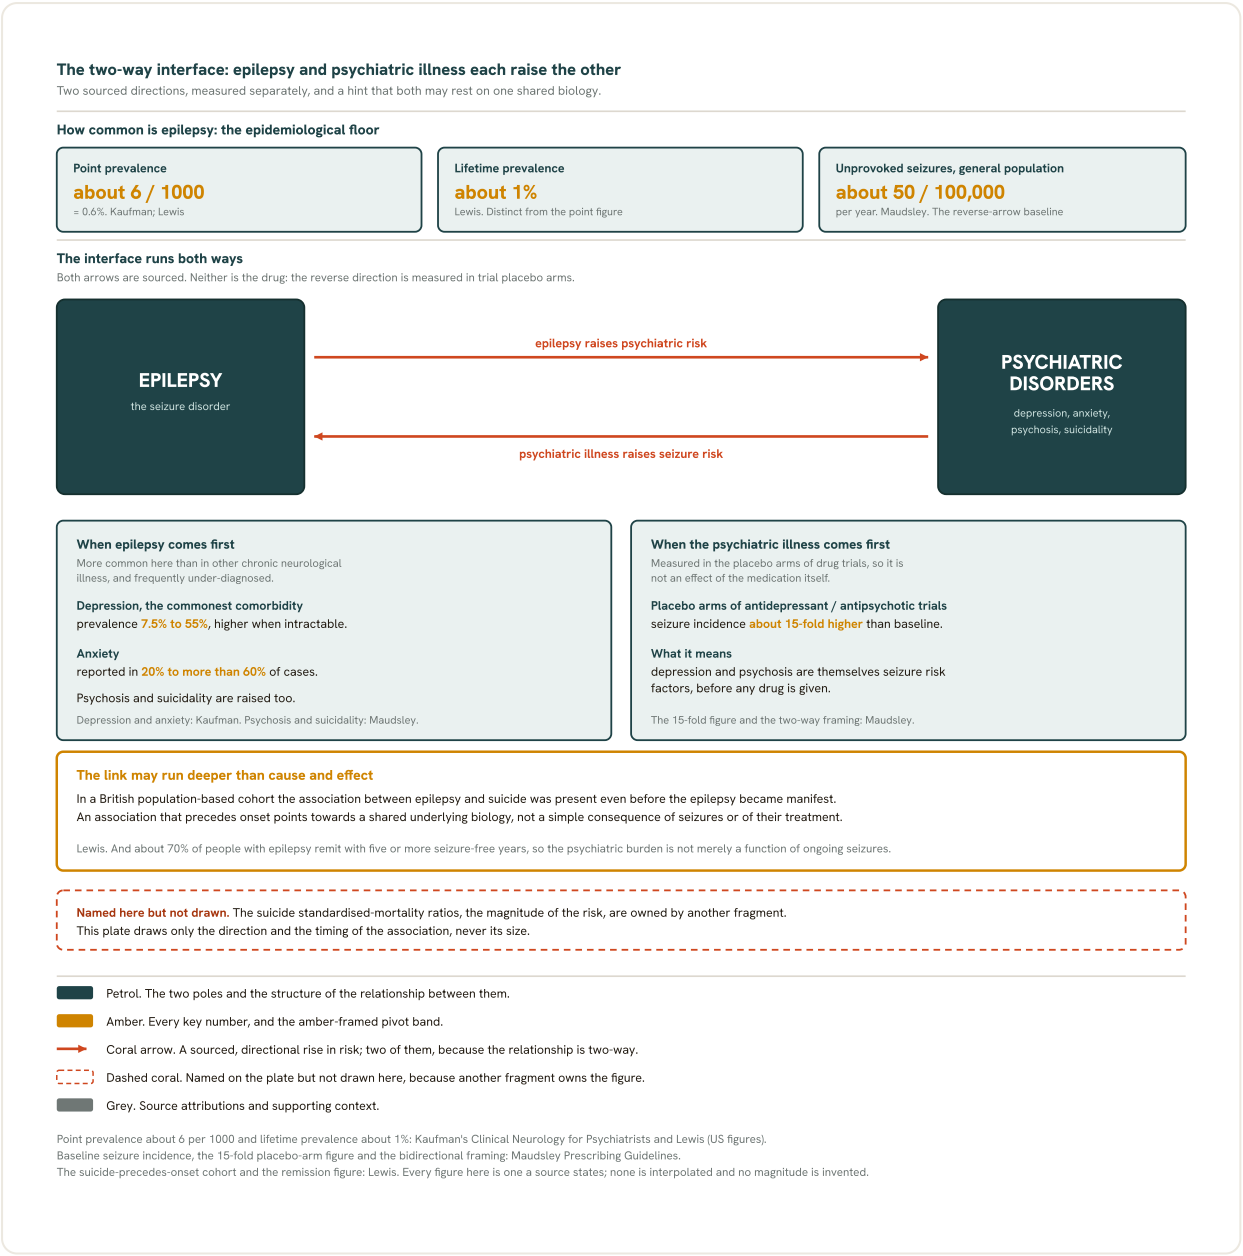


Fig 6.1 **The two-way relationship.** Epilepsy and psychiatric illness set as two poles with a sourced arrow running each way: epilepsy raises the risk of depression (about 7.5% to 55%), anxiety (about 20% to more than 60%), psychosis and suicidality, while psychiatric illness raises seizure risk, shown by a seizure incidence about 15-fold above the general-population baseline of roughly 50 per 100,000 per year in the placebo arms of antidepressant and antipsychotic trials. Because the reverse direction is read from placebo arms, it is a property of the illness rather than of the drug. The amber band carries the deeper point: in a British cohort the epilepsy-suicide association preceded epilepsy onset, which points to a shared biology rather than simple cause and effect. The suicide risk-ratio figures themselves are owned by another fragment and are named but not drawn here.

6.1 How common is epilepsy: prevalence and incidence

Three quantities answer three different questions, and examiners exploit the confusion between them. **Point prevalence** is the proportion of a population who have active epilepsy at a single instant, the figure that tells you how many people in a clinic or on a ward today actually carry the disorder; it sits at about **six people in every 1000**. **Lifetime prevalence** is the proportion who will satisfy the criteria for epilepsy at some time across their lives, and it is necessarily the larger number, of the order of **1%**, because it accumulates everyone who has ever had the condition, including the many now in remission. **Incidence** is different again: it counts new cases arising each year, and for epilepsy it is conventionally expressed as the annual rate of

unprovoked seizures, running at roughly 50 per 100,000 persons in the general population. An unprovoked seizure is one occurring without an immediate provoking cause such as acute metabolic derangement, intoxication or recent head injury, and it is the unprovoked event, rather than any convulsion whatever, that signals the underlying tendency to seize.

Two clarifications keep these numbers honest. A single seizure is not epilepsy; the disorder is defined by an enduring predisposition to recurrent unprovoked seizures, and the formal seizure and epilepsy classification that underpins that definition is the business of the seizure-classification account in this chapter, not of the epidemiology. And acute symptomatic, or provoked, seizures, those driven by a transient insult, are counted separately, because they carry a different prognosis and do not by themselves establish epilepsy. Keeping provoked events out of the count is what makes incidence a measure of the underlying tendency, not of every fit that reaches an emergency department. For the psychiatrist the practical message is twofold: epilepsy is common enough that it will recur throughout a career, in general and liaison settings and in the assessment of paroxysmal behavioural disturbance, and the measures are not interchangeable. Quoting a lifetime figure when asked about point prevalence, or the reverse, signals that a candidate has not grasped what each one measures.

MEASURE	THE QUESTION IT ANSWERS	WHY IT MATTERS TO THE PSYCHIATRIST
Point prevalence	How many people have active epilepsy right now?	Sets the expected caseload burden at any single moment
Lifetime prevalence	How many will carry the diagnosis at some time?	Captures patients now in remission who still bring psychiatric risk and history
Incidence of unprovoked seizures	How many new cases arise each year?	The baseline rate against which seizure risk in psychiatric populations is judged

The three epidemiological quantities distinguished by the question each answers; a lifetime figure and a point figure are not interchangeable, and the incidence of unprovoked seizures is the general-population baseline used later to gauge the reverse direction of the relationship.

6.2 A remitting disorder: why lifetime and point prevalence diverge

The gap between the lifetime figure and the point figure is not a rounding artefact; it is the signature of a condition that most people eventually leave behind. **Remission**, here meaning sustained freedom from seizures over a span of several years whether or not the person stays on medication, is the usual long-term outcome: about **70%** of people with epilepsy eventually reach it. Because the large majority pass into remission, the pool of people with active disease at any instant is smaller than the pool who have ever had it, and that single fact reconciles the two prevalence numbers. It also carries a clinical implication that is easy to miss: a great many patients who carry a past diagnosis of epilepsy are no longer having seizures at all, yet they may still bring the psychiatric and social sequelae of the years in which they were, the interrupted schooling, the surrendered driving licence, the accumulated sense of being marked by the illness.

A minority run a very different course. **Drug-resistant epilepsy**, epilepsy that continues to seize despite adequate trials of appropriate medication, concentrates much of the psychiatric burden, because refractory seizures, repeated hospital contact and the steady erosion of independence are precisely the conditions under which depression and anxiety take hold. The prognostic split therefore matters to the psychiatrist as much as to the neurologist. The person in durable remission and the person with refractory disease sit at opposite ends of a psychiatric risk gradient, and knowing where a given patient falls on it often does more to predict their psychiatric vulnerability than the seizure type alone.

6.3 The bidirectional relationship, in one idea

The organising idea of this whole area is that epilepsy and psychiatric illness are linked both ways. The relationship is **bidirectional**: people with epilepsy carry a raised risk of psychiatric disorder, and people with psychiatric disorder carry a raised risk of going on to develop epilepsy. This is not the older, intuitive picture in which psychiatric symptoms are simply the emotional fallout of living with seizures. That forward pathway is real but it is only one arm of the relationship. The reverse pathway, in which psychiatric illness precedes and predicts seizures, is now equally well established, and it is what lifts the topic from a footnote in a neurology text to a genuine two-way interface deserving a psychiatrist's attention. As the accompanying figure makes clear, the traffic moves in both directions rather than one, and the same psychiatric domains recur on each arm.

MNEMONIC

DAPS

The psychiatric domains for which the two-way link is best demonstrated: **D**epression, **A**nxiety, **P**sychosis and **S**uicidality. Each is more common in people with epilepsy than in the general population, and each, in turn, is associated with a raised risk of seizures.

RULE **The epilepsy-psychiatry relationship is genuinely two-way, not a one-way consequence of seizures.** Any account that explains psychiatric comorbidity solely as a reaction to having epilepsy has described one arm and missed the other. The reverse arm, psychiatric illness as a risk factor for seizures, is where the marks and the clinical surprises both lie.

6.4 The forward arrow: psychiatric disorder is over-represented in epilepsy

That people with epilepsy carry more psychiatric illness than the general population is the older and better-recognised half of the relationship, and the epidemiology is genuinely striking. Depression is the single most common psychiatric **comorbidity** of epilepsy. Reported prevalence spans a wide band, from roughly 7.5% to 55%, with the higher figures clustering among people whose seizures are intractable. The range is broad because studies differ in setting, in how they define a case and in the instruments they use, but even its lower bound sits above the rate seen in most other chronic neurological conditions, and depression here is characteristically both under-recognised and under-treated. The detailed phenomenology and screening of depression in epilepsy belong to the dedicated account of that disorder later in this chapter; what

matters for the epidemiology is its sheer frequency and its claim to first place among the comorbidities.

Anxiety is comorbid almost as often. It is reported in anywhere between about 20% to more than 60% of people with epilepsy, once again with a wide spread that reflects how variably it is defined and measured, and its full clinical treatment is developed in the anxiety-in-epilepsy account of this chapter rather than here. Beyond mood and anxiety, the psychoses of epilepsy and a raised risk of suicide complete the forward picture; each is taken up in full in its own account, and neither should be collapsed into a single number in this topic. The epidemiological point holds across every one of these domains: whatever the precise figure, psychiatric morbidity in epilepsy is not an occasional complication but an expected one.

GOOD TO REMEMBER

Because psychiatric comorbidity in epilepsy is the rule rather than the exception, active enquiry for depression and anxiety belongs in every epilepsy review, not only when the patient volunteers distress. Waiting for a spontaneous complaint is precisely how a frequent, treatable and under-recognised comorbidity is allowed to run unchecked.

PITFALL

Attributing low mood in a person with epilepsy to an understandable reaction to a chronic illness, and therefore watching rather than treating it. The reaction may be genuinely understandable and the depression still be a discrete, biologically driven disorder that warrants the same treatment it would receive in anyone else. Explaining a symptom is not the same as excusing yourself from treating it.

6.5 The reverse arrow: psychiatric illness as a risk factor for seizures

The counterintuitive half of the relationship, and the part examiners most enjoy probing, is that psychiatric illness precedes and predicts epilepsy rather than merely trailing after it. The cleanest quantitative evidence comes from an unlikely place: the placebo arms of randomised trials of antidepressants and antipsychotics. In those arms, where patients receive no active drug at all, the incidence of unprovoked seizures runs about **15-fold higher** than the general-population baseline set out earlier. Since the placebo patients are taking nothing that could lower the seizure threshold, the excess cannot be a drug effect and must instead be attributed to the psychiatric illness itself. Depression and psychosis, on this evidence, behave as seizure risk factors in their own right, and the comparison holds only because there is a stable general-population incidence to measure the trial arms against.

The suicide literature makes the same argument from a different angle. The association between suicide and epilepsy is present even before epilepsy has declared itself: in population-based data the raised risk is already detectable in the years preceding the first seizure. Because that excess cannot be a consequence of seizures that have not yet happened, nor of drugs not yet prescribed, it points instead to a **shared neurobiological diathesis** underlying both conditions. The magnitude of the suicide risk itself, and the specific factors that drive it, are the province of

the suicide-in-epilepsy account and are not quantified here; the lesson at this point is narrower and cleaner, that the timing of the association argues for common ground rather than simple one-way causation.

TRAP **Assuming that seizures in a psychiatric patient must be drug-induced.** Antidepressants and antipsychotics can lower the seizure threshold, but the placebo-arm data show that the underlying illness raises seizure risk on its own, before any drug is given. A candidate who credits every seizure in a depressed or psychotic patient to the medication has quietly denied the reverse arm and lost the point examiners are testing for.

6.6 Reading the two arrows: shared biology, confounding and causation

A two-way association immediately raises the question of what actually connects the two conditions, and here a little epidemiological caution is worth more than a confident answer. Three explanations coexist, and they are not mutually exclusive:

- **Two-way reverse causation** - seizures and the disruption they bring can precipitate mood and anxiety disorders, while depression and psychosis, and the physiological changes that accompany them, can lower the threshold for seizures.
- **Confounding by shared antecedents** - early adversity, head injury, substance use and social disadvantage each raise the risk of both epilepsy and psychiatric illness, and can manufacture an association without either condition causing the other.
- **A common biological substrate** - a shared vulnerability expressed as both a tendency to seize and a tendency to psychiatric symptoms.

What that substrate might actually be belongs elsewhere in this chapter. Its detailed account is the business of the chapter's topic on kindling and the limbic basis of behaviour, and the full anatomy of the structures involved is set out in the Functional Neuroanatomy and Neurophysiology notes; both are named here rather than re-taught. The discipline for the reader is to resist collapsing the three explanations into one. The two-way finding does not, on its own, prove shared biology, since it is equally compatible with reciprocal causation and with **confounding**. What tips the balance towards a common substrate is the observation, from the suicide data, that risk is already raised before epilepsy begins, a pattern that seizures and their treatment simply cannot explain. It is this temporal ordering, more than the bare fact of association, that carries the argument. Applied with that care rather than as a slogan, it is **reverse causation** reasoning at its most defensible.

6.7 Why the epidemiology changes practice

Numbers earn their place only when they change behaviour. The prevalence and the two-way findings converge on a small set of habits that separate careful from careless practice across the neurology-psychiatry border:

- Screen actively for depression and anxiety in everyone with epilepsy, given how frequently they occur and how quietly they are missed when the clinician waits for a spontaneous complaint.

- Take a seizure and a spell history in psychiatric patients, remembering that the illness itself raises seizure risk and that a first seizure can surface on that background rather than only as a drug effect.
- Treat comorbid psychiatric disorder on its own merits rather than writing it off as an inevitable reaction, while staying alert to seizure-threshold effects when selecting an agent.

The last point sits at the centre of the therapeutic tension that runs through the rest of this chapter. Psychiatric illness in epilepsy must be treated, yet the drugs that treat it can influence seizure control, and the drugs that control seizures can in turn influence mood, cognition and behaviour. The epidemiology is exactly what makes that tension unavoidable: if comorbidity were rare, the prescribing dilemma would seldom arise, but because it is the norm, every clinician working at this interface will meet it, and meet it often. The specifics of choosing drugs against the seizure threshold, and of the mood and cognitive effects of the antiepileptics themselves, belong to later topics; the point to carry from here is that the frequency of comorbidity is what forces those choices into the open.

IN INDIA

The population burden of epilepsy is amplified by a wide treatment gap, with a large share of people who have active seizures receiving no consistent antiepileptic care, particularly in rural and under-served districts. In the same high-volume clinics, where a consultation is understandably consumed by seizures and drug supply, the psychiatric comorbidity this topic insists on is even easier to overlook, which is exactly where deliberate, structured enquiry earns its place.

6.8 Pulling the picture together

Common, chronic, mostly remitting, and bound to psychiatric illness in both directions. The topic resolves into a short chain. Epilepsy is common, with a point prevalence lower than its lifetime prevalence because most people eventually remit, and the gap between those two figures is the natural history of the disease made visible rather than a contradiction to be explained away. The relationship with psychiatric disorder is genuinely two-way: the forward arm, in which depression leads a frequent and under-recognised set of comorbidities, and the reverse arm, in which psychiatric illness raises seizure risk in its own right, shown most cleanly by the excess seizures in trial placebo arms and by a suicide association that predates the epilepsy itself. That last point nudges the explanation towards a shared substrate without proving it. Everything the rest of the chapter does, from the peri-ictal syndromes to the prescribing dilemmas, stands on this frame.

~6 / 1000	~1%	~15-fold higher	~70%
POINT PREVALENCE	LIFETIME PREVALENCE	SEIZURES, PLACEBO ARMS VS POPULATION	EVENTUALLY REMIT

Epilepsy is common and mostly remitting, and its relationship with psychiatric illness is bidirectional: seizures raise the risk of psychiatric disorder, and psychiatric disorder raises the risk of seizures, a two-way link best demonstrated for depression, anxiety, psychosis and suicidality.

7 · EEG fundamentals and ictal versus interictal discharges in psychiatric assessment

The electroencephalogram is the one bedside investigation that speaks directly to the border between epilepsy and psychiatry. It is also the investigation psychiatrists most often misread, trusting a normal tracing too far or over-reading an incidental sharp transient in a patient whose real trouble is functional rather than epileptic. Electroencephalography was born inside psychiatry, and the discipline has kept a fascination with the electrical brain ever since. A psychiatrist rarely reports an EEG but is constantly asked to act on one: to decide whether an odd spell was a seizure, whether an atypical psychosis sits on an epileptic substrate, or whether a normal study can be allowed to close a question. This section builds the vocabulary needed to read a report critically rather than take its conclusion on trust. It covers the normal background rhythms and their frequency bands, the morphology of the epileptiform discharge, the sharp limits on what the scalp can and cannot see, and above all the distinction between activity recorded between seizures and activity recorded during one. That last contrast, interictal against ictal, is the electrographic spine of the whole framework that later classifies psychiatric syndromes by their timing to the seizure.



Interictal EEG: Javidan M, Epilepsy Res Treat 2012;2012:637430, Fig 2, CC BY 3.0



Ictal EEG: Javidan M, Epilepsy Res Treat 2012;2012:637430, Fig 3, CC BY 3.0

Fig 7.1 Interictal versus ictal EEG in temporal lobe epilepsy. Above, the record between seizures: intermittent sharp waves over the anterior and sphenoidal temporal channels, the electrical signature that points to the epileptogenic focus without any clinical event. Below, the record during a seizure: a rhythmic discharge builds and evolves over the temporal channels, the hallmark of a focal seizure in progress. The distinction is the whole point of the recording. Interictal sharpening localises the focus; the evolving ictal rhythm confirms a seizure, and the two must never be read as the same thing.

7.1 Where the EEG came from and what the scalp records

A psychiatrist's instrument from the start. The recording of rhythmic electrical activity from the intact human scalp was first achieved by **Hans Berger**, a German psychiatrist, who published the human electroencephalogram in 1929 and in doing so founded electroencephalography as a discipline. The rhythm he described over the back of the head in relaxed wakefulness became the reference point for everything that followed. What the surface electrode actually measures is worth stating plainly, because most misreadings begin with a wrong mental model of the signal.

The EEG does not record action potentials or the firing of single neurons. It records the summed excitatory and inhibitory postsynaptic potentials of large populations of cortical pyramidal cells, whose parallel radial orientation lets their minute extracellular currents add together into a field large enough to reach the scalp. For that summation to become visible, thousands of neurons must be active in near synchrony; sparse or asynchronous activity cancels out and disappears. The signal is then further attenuated and smeared by the meninges, skull and scalp that lie between cortex and electrode, so that what reaches the recording is a blurred, low-amplitude average of a great deal of cortical work. This is why the EEG is exquisitely sensitive to changes in cortical state, the very thing that shifts in delirium, encephalopathy and sedation, yet spatially coarse and blind to much of what the brain is doing. Holding that model steady, a synchronised population average heavily filtered by the head, makes sense of both the power and the limits that the rest of this topic explores.

7.2 The frequency bands of the normal background

Four bands, slowest to fastest. The continuous background of the EEG is described by carving its frequencies into named bands, and four of them carry the clinical weight a psychiatrist needs. Following the conventions of the standard primer How to Read an EEG, they are set out in the accompanying table. The **posterior dominant rhythm** of the healthy, relaxed adult with eyes closed sits in the alpha range and attenuates when the eyes open or attention engages; faster beta activity accompanies alertness and is a familiar signature of sedative-hypnotic drugs; the slower theta and delta frequencies are physiological in drowsiness, sleep and childhood but become pathological when they intrude on the waking adult record. That intrusion is the single most common reason a psychiatrist meets an abnormal EEG. Diffuse slowing of the background into the theta and delta range is the electrographic hallmark of a diffuse encephalopathy, the state that underlies delirium, whereas slowing confined to one region points to a focal structural lesion beneath it. One caution belongs with the numbers. The exact cut-offs between adjacent bands, and the theta boundary in particular, are drawn slightly differently by different authorities, so a boundary figure should be held as a working convention rather than a law of nature.

BAND	FREQUENCY	WHERE IT BELONGS
Delta	0.5 to 4 Hz	Deep sleep in health; pathological as diffuse or focal slowing in the waking adult
Theta	4 to 7 Hz	Drowsiness and light sleep; normal and prominent in childhood
Alpha	8 to 13 Hz	The posterior dominant rhythm of relaxed, eyes-closed wakefulness
Beta	14 to 30 Hz	Active concentration; accentuated by sedative-hypnotic drugs

The four clinically weighted frequency bands of the background, from slowest to fastest.

7.3 What the scalp can and cannot see

Most of the cortex is effectively silent to the scalp. The most consequential limitation of the routine study is spatial. Only about **one-third** of the cortical surface, the crowns of the superficial gyri lying just beneath the electrodes, generates a signal the scalp can read well. The remainder is electrographically quiet from the surface: the cortex buried in the depths of the sulci, the basal surfaces, the mesial walls between the hemispheres and the insula folded deep within the Sylvian fissure all sit too far from any electrode, or are oriented so that their fields point away from it, to register reliably. This single fact of geometry drives much of the clinical trouble a psychiatrist inherits, because the regions that are hardest to see are precisely the ones of greatest psychiatric interest.

- The mesial temporal structures, the seat of temporal lobe epilepsy, lie deep and generate weak, often absent, scalp signals.
- The mesial and orbital frontal surfaces, the substrate of frontal lobe epilepsy, are similarly buried and poorly represented.
- Deep and basal sources in general are attenuated by distance and by the intervening tissue before they ever reach the recording.

The clinical reading of this is uncomfortable but essential: a scalp EEG can be repeatedly and genuinely normal in a patient who is having unmistakable seizures from a deep or mesial focus. The tracing has not failed; it is reporting faithfully on the third of the cortex it can reach and staying silent about the rest. For the psychiatrist weighing an atypical presentation that might be epileptic, this means a normal surface study can never be treated as a clean bill of health for the deep structures where the epilepsy-behaviour interface actually lives.

TRAP **Reading a normal interictal EEG as proof the patient does not have epilepsy.** The scalp sees only the convexity, and a single routine tracing captures a minority of established cases; a mesial temporal or deep frontal focus can discharge freely while the surface record stays stubbornly normal. A normal EEG lowers the probability of epilepsy, it never closes the question.

7.4 Epileptiform discharges: the spike and the sharp wave

Two pointed transients that differ only in duration. When the EEG is abnormal in a way that suggests epilepsy specifically, it is usually because of an **epileptiform discharge**: a brief, pointed waveform that stands out abruptly from the surrounding background, typically followed by a slower after-going wave. Two forms of this transient are distinguished, and the difference between them is purely one of duration. A **spike** is the briefer of the two, lasting **longer than 20 but under 70 milliseconds**, while a **sharp wave** is the broader form, lasting **between 70 and 200 milliseconds**. Both are pointed, both arise from the synchronised, paroxysmal depolarisation of a population of neurons, and both belong to the family of interictal epileptiform abnormalities; morphology does not separate them, only the clock does. It matters to hold this firmly, because not every sharp-looking transient in an EEG is epileptiform. The background is full of benign sharp elements and artefacts that a careful electroencephalographer distinguishes from a true discharge by their field, their after-going slow wave, their disruption of

the surrounding rhythm and their location, and the untrained eye that calls every pointed waveform a spike will generate false alarms. The genuine epileptiform discharge is a specific object with defined temporal limits, not merely anything that happens to look sharp.

MNEMONIC

Spike is Speedy, Sharp is Slower

The two interictal transients are separated on one axis only, their duration: the **sp**ike is the briefer, faster peak and the **sh**arp wave is the broader, longer one. Both are pointed epileptiform discharges, so it is the clock, and not the shape, that tells them apart.

7.5 Reading the discharge: probability, not proof

An epileptiform discharge is a strong pointer, not a verdict. The single most important interpretive principle is that a spike or a sharp wave is not pathognomonic of epilepsy. It is not, on its own, proof that the person has the disorder; what it establishes is a strong predisposition to seizures rather than the disorder itself, and a small proportion of entirely healthy people carry such discharges without ever having a seizure. The way to make this quantitative is to separate two properties that candidates routinely confuse. The epileptiform EEG is highly specific but poorly sensitive for epilepsy, and that asymmetry is the whole point. Its high specificity is what gives a discharge its weight: when one is genuinely present, epilepsy becomes a strong bet. Its low sensitivity is the other blade: most people with epilepsy will not show a discharge on any given routine recording, so the absence of one proves very little. The prevalence of epileptiform discharges in the healthy population anchors the specificity from the other side, low enough that an incidental discharge is uncommon yet not so low that it can be dismissed out of hand. Reading a report therefore means holding the result against the clinical pre-test probability. In a patient with a convincing seizure history, a discharge confirms what was already likely; in a patient with none, the same discharge should prompt caution rather than a diagnosis, because the base rate of epilepsy in that person is low and even a highly specific test throws the occasional false positive.

~95%	~30%	<1%	<4%
SPECIFIC FOR EPILEPSY	SENSITIVE FOR EPILEPSY	OF HEALTHY ADULTS	OF HEALTHY CHILDREN

RULE **The epileptiform EEG confirms far more reliably than it excludes.** A discharge is a strong pointer because the finding is highly specific, but its poor sensitivity means a single recording misses most people who have epilepsy. Read every result against the clinical pre-test probability: the EEG moves the odds, it does not settle the diagnosis by itself.

PITFALL

The mirror error to the examination trap is over-reading. A psychiatrist who orders an EEG as a routine screen for organicity in every atypical presentation will sooner or later be handed a report noting a sharp transient, and will be tempted to rediagnose on the strength of it. Because such discharges occur in a small minority of entirely healthy people, an incidental finding in someone with no clinical seizures rarely changes the picture and should not, by itself, convert a functional or affective presentation into epilepsy. Order the study to answer a defined question, not as a fishing expedition.

7.6 Why one tracing is rarely enough

The single routine EEG is the least sensitive form of the test. The low sensitivity of the interictal recording has a direct practical consequence: **fewer than half** of people with established epilepsy show any epileptiform abnormality on a single routine scalp EEG. A normal first study is therefore the expected result in a large share of genuine cases, and it must never be allowed to overturn a good clinical history. The corollary is that yield can be raised, and the ways of raising it are the practical heart of ordering the test well. Timing is the most powerful lever, since the chance of capturing a discharge is highest soon after a seizure and greatest within roughly the first 24 hours of an event, so an EEG requested promptly after a first seizure earns far more than one deferred for weeks. Beyond timing, the sensitivity of the study itself can be increased.

- Repeating the recording, since serial studies accumulate the chance of catching an intermittent discharge that any one tracing misses.
- Recording during drowsiness and sleep, and using sleep deprivation, which brings out discharges that a fully alert cortex suppresses.
- Prolonging the recording, since a longer or ambulatory study simply samples more time and more brain states.

GOOD TO REMEMBER

Two habits improve the yield without any new equipment. First, record early, because an EEG requested at once after a first event is worth far more than one booked for a month later. Second, read a single normal routine study for exactly what it is, the least sensitive form of the test, rather than as evidence that epilepsy has been excluded. Between them, these two habits prevent the commonest ways a psychiatrist is misled by a scalp recording.

IN INDIA

A routine scalp EEG is inexpensive and widely available across the country, and that cuts both ways. Because it is so easy to obtain, it is frequently over-requested in psychiatric and functional-seizure workups and its incidental findings over-interpreted, while the studies that actually raise the yield, prolonged, sleep-deprived and video-EEG telemetry, remain concentrated in tertiary and academic centres. The practical stance for a psychiatrist in a district setting is to lean on a careful clinical history, to treat a single normal routine EEG as non-exclusory rather than reassuring, and to refer onward when the question genuinely needs activity captured during an event, including the separation of functional from epileptic seizures.

7.7 Interictal against ictal: the distinction that organises everything

The contrast the whole peri-ictal framework turns on. Everything discussed so far, the spike and the sharp wave, has been an **interictal discharge**: an abnormality recorded between seizures, when the patient is clinically well. It testifies to a brain that is capable of seizing, but it is not itself a seizure. An **ictal pattern**, by contrast, is the electrographic seizure itself: not a single transient but a rhythmic, evolving run of activity that builds, changes in frequency and amplitude, spreads across the head and then subsides. The two were mapped onto epilepsy early in the field's history. **Gibbs and Jasper** first reported the interictal discharge of focal epilepsy in 1936, and the most famous ictal and interictal signature of all, the generalised **3 Hz** spike-and-wave discharge of idiopathic generalised epilepsy with absences, was demonstrated by **Gibbs, Davis and Lennox** in 1935. The distinction is not a piece of electrophysiological trivia; it is the difference between a finding that suggests a diathesis and a finding that catches the disorder in the act.

FEATURE	INTERICTAL DISCHARGE	ICTAL PATTERN
When it occurs	Between seizures, with the patient clinically well	During the seizure itself
Form	A brief, stereotyped transient, the spike or sharp wave	A rhythmic run that evolves in frequency, amplitude and distribution
What it establishes	A predisposition; a brain able to seize	That a seizure is happening at that moment
Bearing on behaviour	Supports an epileptic tendency but cannot time a given spell	If locked to an observed spell, proves the spell is that seizure

The interictal and ictal patterns compared, and what each does and does not prove.

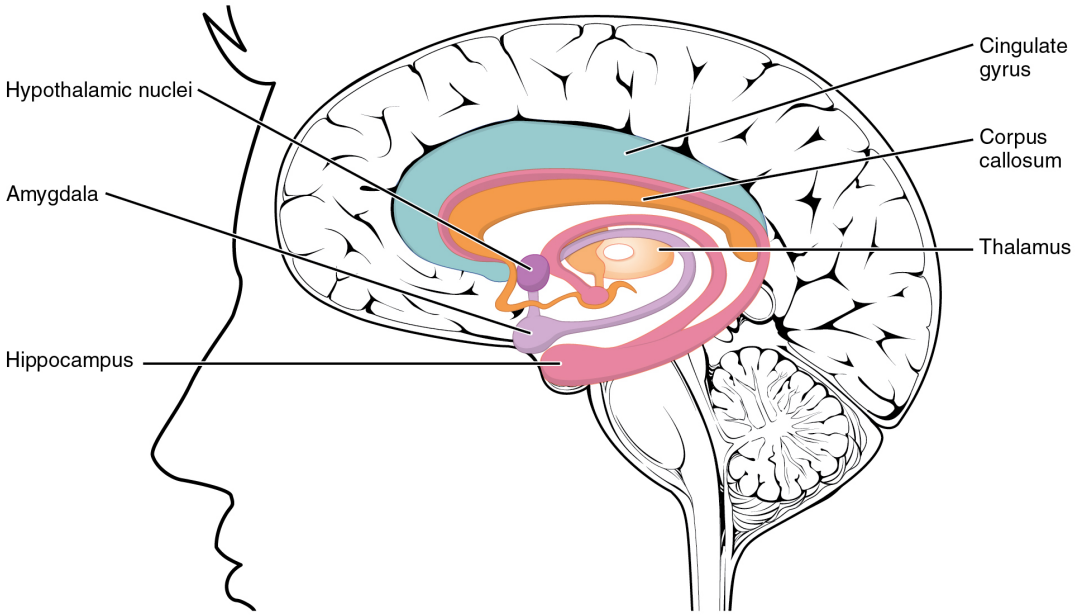
The clinical payoff of the distinction is precise. An interictal discharge supports the presence of an epileptic tendency but cannot, on its own, tell you that any particular episode was a seizure. An ictal pattern does exactly that: when a rhythmic electrographic seizure is recorded in lockstep with an observed behavioural change, a spell of fear, a period of altered awareness, an automatism, it proves that the behaviour in front of you is that seizure and not something else. This is why the capture of an event on a recording is the aim of the more elaborate studies, and why the interictal-ictal distinction is the electrographic foundation of the framework that

classifies psychiatric syndromes by their timing to the seizure. To place a preictal, ictal, postictal or interictal label on a behavioural disturbance is, at bottom, to make a statement about its relationship to an electrographic seizure, and that statement rests on the difference this section has drawn.

> Interictal discharges tell you a brain is capable of seizing; an ictal pattern tells you it is seizing at that moment, and only the second can prove that the behaviour in front of you is itself a seizure.

8 · Limbic system, kindling and the neurobiology linking epilepsy to behaviour

Epilepsy is at root a disorder of abnormal electrical discharge, yet the reason it earns a place in a psychiatry examination is that it so reliably changes behaviour, mood and thought. The question a candidate is really being asked, whenever this material appears, is a mechanistic one: by what route does a paroxysm of electrical activity in the brain translate into a durable psychiatric change that can outlast any single seizure by years. This section is the chapter's mechanism step. It does not catalogue the peri-ictal syndromes, the personality change of long-standing epilepsy or the several psychoses of epilepsy, each of which is described in its own right elsewhere in this account; it supplies the neurobiology that makes all of them intelligible. Two ideas carry the weight. The first is anatomical: the seizures that matter most for behaviour arise in or spread to the limbic system, the emotional and mnemonic core of the brain. The second is physiological and is the heart of the topic: the phenomenon of kindling, a laboratory model in which repeated limbic seizure activity leaves a lasting change in how the brain works. Hold those two together, a vulnerable anatomical stage and a mechanism that makes seizure activity self-perpetuating and behaviourally consequential, and the whole clinical picture of epilepsy-related psychiatric disorder acquires a spine.



OpenStax College, Anatomy and Physiology, "The Limbic Lobe", CC BY 3.0, via Wikimedia Commons

Fig 8.1 The limbic circuit and the mesial temporal structures that matter in temporal lobe epilepsy. The amygdala and the hippocampus sit at the medial edge of the temporal lobe and are the commonest origin of focal seizures that carry psychiatric colour; through the limbic ring they connect to the cingulate gyrus, the hypothalamus and the thalamus. This is the anatomy behind the affective and experiential aura, behind the memory disturbance of mesial temporal sclerosis, and behind the kindling model in which repeated limbic discharge lowers the threshold for the next.

8.1 The limbic system as the anatomical stage

Why behaviour and seizures meet in the same territory. The limbic structures are where the epilepsy-behaviour interface physically lives, and this is no coincidence. The same medial temporal and basal forebrain regions that generate and propagate the seizures of greatest psychiatric interest are the regions that subserve emotion, motivation, fear and the laying down of memory. When a discharge arises in or recruits this circuitry, it is not activating neutral tissue; it is driving the very apparatus of affect and salience, which is why limbic seizures so often carry an emotional or experiential colour and why chronic limbic epilepsy has behavioural sequelae that a purely motor epilepsy does not. The limbic system also has a low threshold for the propagation and self-reinforcement of abnormal activity, a property that becomes central once kindling is introduced. The full anatomy of this system, its component structures, their interconnecting circuitry and its wider projections, is a neuroscience topic in its own right and is set out in the Functional Neuroanatomy and Neurophysiology notes; it is named here and not re-taught. The one name to carry across is the classical loop itself, the **Papez circuit**, which runs from the hippocampus through the fornix to the mammillary bodies, on to the anterior thalamus and the cingulate gyrus and back by way of the entorhinal cortex, the anatomical ring through which emotion and memory are bound together; the amygdala sits within the limbic system but outside that canonical loop. What matters for the present purpose is only the functional consequence: the limbic system is simultaneously the brain's principal seizure-prone gateway and its emotional engine, so activity that disturbs it disturbs behaviour almost by definition.

Everything that follows is an account of how repeated disturbance of this stage produces changes that persist.

8.2 Kindling: the experimental paradigm

A small, repeated stimulus that grows its own seizure. The single most important laboratory model linking seizure activity to lasting change is **kindling**. In the classic preparation, a brief, low-intensity electrical stimulus is delivered periodically to the amygdala or another limbic structure. The first stimulation is deliberately weak: it evokes only a brief electrographic seizure, an **afterdischarge** visible on the recording, with no outward change in the animal's behaviour at all. Left there, nothing would come of it. The defining observation is what happens on repetition. When the same modest stimulus is applied again and again over successive sessions, the evoked response does not stay constant; it escalates. The afterdischarges lengthen and spread, and behavioural accompaniments appear and intensify in a stereotyped progression, until a stimulus that once produced only a silent electrographic flicker now reliably provokes a full, generalised behavioural seizure. As the standard pharmacology account in Goodman and Gilman's therapeutics text describes the model, it takes only **about 10 to 20** such stimulations for this progressive intensification to unfold. The power of the paradigm is precisely that the stimulus never changes. Nothing is made stronger from the outside; the brain itself becomes progressively more responsive to an unchanging input, manufacturing an ever-larger seizure from an ever-constant spark. That is the observation the rest of the topic has to explain, and it is the reason kindling, rather than any single acute seizure, is treated as the model of how epilepsy leaves a mark.

- An initial weak stimulus evokes only a brief afterdischarge, with no behavioural change.
- Repetition of the identical stimulus lengthens and spreads the electrographic seizure.
- Behavioural seizure components appear and escalate in a stereotyped sequence.
- The endpoint is a full generalised seizure elicited by the same stimulus that once did nothing.

8.3 Why the amygdala kindles most readily

The emotional structure is also the most excitable one. Kindling can be established at several limbic sites, but they are not equivalent, and the difference is itself instructive. The **amygdala** is the ideal site for kindling because seizure activity can be induced there most rapidly, a point emphasised in the account of the phenomenon in Panksepp's Affective Neuroscience. This is not an incidental piece of laboratory trivia. The amygdala is the brain's principal mediator of fear, threat appraisal and the emotional tagging of experience, so the fact that it is also the tissue most readily driven into a self-amplifying seizure is exactly the anatomical coincidence that makes limbic epilepsy behaviourally loud. A structure that kindles quickly is a structure whose repeated activation most efficiently installs the lasting change that kindling produces, and when that structure is the seat of emotion, the lasting change is liable to be emotional and behavioural in kind. The rapidity of amygdalar kindling therefore does double duty in the argument: it identifies the most vulnerable node in the limbic circuit, and it explains why the psychiatric consequences of chronic limbic seizures so often take the form of altered affect, fearfulness and reactivity

rather than, say, a pure motor or sensory residue. The anatomy of the amygdala as a structure belongs to the neuroanatomy account; what is owned here is the physiological claim that it is the fastest-kindling site and the behavioural inference that follows from it.

8.4 The nature of the change: functional, not structural

A lasting reorganisation you cannot see on a scan. The most consequential feature of the kindled state is what kind of change it is. Kindling induces a permanent change in the functional organisation of the brain with no change evident at the structural level. Once an animal is kindled it stays kindled: the heightened seizure susceptibility does not fade with the passage of time, yet nothing corresponding can be found when the tissue is examined for gross structural damage. The change is in how the circuit works, in the strength and pattern of its synaptic connections and the ease with which activity propagates through it, not in any visible lesion. Conceptually this belongs with the broad family of use-dependent plasticity, the same class of lasting, activity-driven strengthening of synaptic transmission that underlies ordinary learning and memory; kindling can be read as a pathological cousin of that machinery, in which the thing being learned is the seizure itself. This is the property that makes kindling such a satisfying bridge between an electrical event and an enduring behavioural one. A single seizure comes and goes; a functional reorganisation that persists indefinitely is exactly the kind of durable substrate a durable behavioural change requires. It also disciplines how the clinician thinks about a normal scan in a patient whose behaviour has shifted after years of seizures, because the model predicts precisely the situation of a real change in brain function with an unremarkable image.

MNEMONIC

Small spark, lasting fire

Kindling's lesson compressed: a stimulus too **small** to change behaviour at first, once repeated, leaves a **lasting** reorganisation of function that does not fade and does not show on structural imaging. The name is the memory aid - a fire kindled from small sparks that, once lit, stays lit.

PITFALL

A normal MRI is taken to mean the brain is structurally intact and therefore that a patient's behavioural or affective change cannot be seizure-related. Kindling predicts the opposite situation directly: a genuine, permanent reorganisation of function can coexist with a wholly unremarkable structural scan. Reassurance drawn from clean imaging should not be allowed to dismiss a real, biologically grounded change in a patient with long-standing limbic epilepsy.

8.5 One phenomenon reached by several routes

Not an artefact of the electrode. A reasonable first suspicion is that kindling is a quirk of pushing current through an electrode, an experimental artefact with no bearing on natural disease. The convergence of routes is what answers that objection. Kindling can also be induced by periodically administered seizure-inducing drugs, so-called **chemical kindling**, and by repeated loud auditory stimulation in susceptible animals, or **auditory kindling**. That the same progressive, self-amplifying, permanent phenomenon can be produced electrically,

pharmacologically and by a purely sensory stimulus tells us it is a property of the nervous tissue and not of any one method of provoking it. Whatever repeatedly drives the limbic circuitry into brief seizures, whether an implanted electrode, a convulsant compound or an intense recurring sound, sets the same machinery in motion. This generality is the reason kindling can be offered as a model of human disease at all. Human beings are not stimulated through amygdalar electrodes, but they do sustain recurrent limbic seizures from natural causes over years, and the multiplicity of experimental routes makes it plausible that the natural process engages the identical mechanism. The routes differ; the outcome, a lasting escalation of seizure susceptibility, does not.

- Electrical: brief, low-intensity periodic stimulation of a limbic site, the original paradigm.
- Chemical: periodically administered seizure-inducing drugs producing the same escalation.
- Auditory: repeated loud auditory stimulation in susceptible animals.

8.6 From kindled seizure to altered behaviour

This is the step the whole topic exists to make. A model of escalating seizures is only half of what a psychiatrist needs; the other half is the link from seizure activity to behaviour. Kindling supplies it. Repeated kindling produces a durable change in behaviour in animals, and it is this observation, not any single acute seizure, that provides the neurobiological model linking limbic seizure activity to altered behaviour. The animal does not merely become more seizure-prone; it becomes, in lasting ways, a behaviourally different animal. That is the bridge in its cleanest experimental form. From it, two lines of clinical thought are drawn, and the discipline of this section lies in naming them without annexing them. The enduring behavioural residue of chronic human limbic epilepsy, the interictal personality change of long-standing temporal lobe epilepsy, is the human counterpart of this animal observation, but it is a topic in its own right and is not taught here. Separately, the concept of kindling is invoked as an etiological model to help explain the emergence of psychotic phenomena in epilepsy, offering a mechanism by which a psychiatric syndrome can arise gradually over years of accumulating limbic seizure activity rather than appearing fully formed. The psychoses of epilepsy themselves, ictal, postictal and chronic interictal, are described elsewhere; kindling is only the candidate mechanism that makes their slow emergence intelligible. What this fragment owns is the bridge; what the bridge leads to is owned by others.

RULE **Repeated limbic seizure activity leaves a permanent functional change that durably alters behaviour.** This single proposition is the neurobiological link between epilepsy and its psychiatric consequences. It is why a chronic epilepsy can reshape affect and conduct over years, and why the relevant lesion is one of function rather than structure. Every peri-ictal syndrome, personality change and epileptic psychosis discussed elsewhere ultimately draws on this mechanism.

GOOD TO REMEMBER

When a patient with years of limbic epilepsy shows a settled change in mood, reactivity or thought, the kindling model gives you a reason to treat that change as biologically grounded rather than as a purely psychological reaction to having a chronic illness. It does not license a deterministic story, but it does argue against dismissing a durable behavioural shift as merely reactive in someone whose limbic circuitry has been seizing for a long time.

8.7 Epileptogenesis: how an epileptic focus is acquired

From a single insult to a chronic disorder. Kindling models the progressive facilitation of seizures; a related but distinct body of work models how a brain that has never been epileptic acquires a chronic epilepsy, the process of **epileptogenesis**. Here the experimental insult is not a train of small stimulations but a single severe one. In the status-epilepticus models, a chemoconvulsant, either **kainic acid** or **pilocarpine**, is used to provoke prolonged seizure activity, and the animal is then followed. After a latent interval of weeks, during which the animal appears well, spontaneous recurrent seizures emerge and the hippocampus is found to be sclerosed. The sequence is the point: a single acute insult, a silent maturation period, and then a self-sustaining epilepsy with a characteristic hippocampal lesion. This maps with striking directness onto a well-recognised human pathway, in which complicated febrile seizures in early childhood precede, after a long silent interval, the mesial temporal sclerosis that underlies temporal lobe epilepsy, the substrate detailed in the dedicated account of that condition. The animal work makes that clinical association mechanistically credible: it shows, in a controlled setting, that an early episode of intense limbic seizure activity can causally produce the sclerotic focus years before the chronic epilepsy declares itself. The accompanying figure sets kindling and this chemoconvulsant model side by side as the two complementary halves of the neurobiological story, one explaining escalation, the other explaining acquisition.

FEATURE	KINDLING	CHEMOCONVULSANT STATUS MODEL
Provoking insult	Brief, low-intensity stimulation, repeated periodically	A single episode of chemically induced prolonged seizure activity
What it models	Progressive facilitation of seizures and durable behavioural change	Acquisition of a chronic epilepsy from an initial insult
Time course	Stepwise intensification across successive stimulations	A silent latent interval, then spontaneous seizures
End state	Permanent functional reorganisation, no structural lesion	Spontaneous recurrent seizures with hippocampal sclerosis
Human parallel	Limbic seizures linked to lasting altered behaviour	Complicated febrile seizures preceding mesial temporal sclerosis

The two complementary limbic models: kindling explains escalation and behavioural change, the chemoconvulsant status model explains how a focus is acquired.

8.8 What the model establishes, and what it does not

A powerful bridge, held at the right altitude. The value of kindling to the psychiatrist is that it converts a plausible clinical intuition, that years of limbic seizures can change a person, into a demonstrated biological mechanism with a defined form: escalation from an unchanging stimulus, a permanent functional reorganisation, and a durable behavioural residue, reached by several independent routes. That is a great deal, and it is what the marks reward. What the model does not do is prove, in any individual human patient, that a given psychiatric syndrome was caused by kindling. It is a model and a heuristic, established in animals and invoked to make human phenomena intelligible, not a diagnostic test that can be applied at the bedside. The distance between a controlled animal preparation and a particular person's psychosis or personality change is real and must be respected. Read at the right altitude, kindling earns its central place: it is the reason the chapter can speak of an epilepsy-behaviour interface at all, the mechanism that underwrites the peri-ictal syndromes, the personality change and the epileptic psychoses without itself being any of them.

Amygdala	Functional	Durable	A model
THE FASTEST-KINDLING SITE	PERMANENT REORGANISATION, NOT STRUCTURAL	BEHAVIOURAL CHANGE IN ANIMALS	FOR EMERGENT PSYCHOSIS, NOT A PROOF

TRAP **Treating kindling as proven causation in the individual patient rather than as an animal model and heuristic.** Candidates who have grasped the mechanism are tempted to state that a patient's epilepsy caused their psychosis or personality change by kindling. The evidence supports kindling as a model that explains how such changes can emerge; it does not license a causal verdict in a single case. State the mechanism as a model, and the direction of the claim survives.

Kindling shows that repeated limbic seizure activity leaves a permanent, purely functional reorganisation of the brain that durably alters behaviour, and that single fact is the neurobiological bridge from an electrical disorder to its psychiatric consequences.

The peri-ictal axis: syndromes by their timing to the seizure

9 · Classification of peri-ictal psychiatric syndromes by temporal relationship to seizure

The single most useful question at the bedside is not what the symptom is, but when it sits. A patient with epilepsy who turns frightened, confused, elated, paranoid or profoundly low presents the psychiatrist with a phenomenology that, taken in cross-section, can be indistinguishable from a primary psychiatric disorder. What tells them apart, and what organises the otherwise sprawling behavioural landscape of epilepsy into something teachable, is the relationship each disturbance holds to the seizure itself. That temporal relationship is the classifying axis of this topic, and mastering it earns disproportionate marks because it is the scaffold on which every other peri-ictal syndrome hangs. Get the timing right and a wide

differential collapses to a short candidate list; get it wrong and a self-limiting, seizure-driven state is mistaken for a lifelong illness.

The anchor of the whole scheme is the seizure event itself, against which every psychiatric phenomenon is dated. Working outward from that anchor, a disturbance may coincide with the discharge, may cluster in the hours around it, may occupy the long stretch between attacks, or may float free of any dependable temporal tie. Each position carries its own diagnostic weight, its own list of candidate syndromes and, decisively, its own management logic. This account builds the axis from its vocabulary upward, places the categories of psychiatric disturbance onto it, and then draws out the one management principle the timing dictates. The syndromes that occupy each position are developed in their own right elsewhere in this chapter; here they are only located, never re-taught.

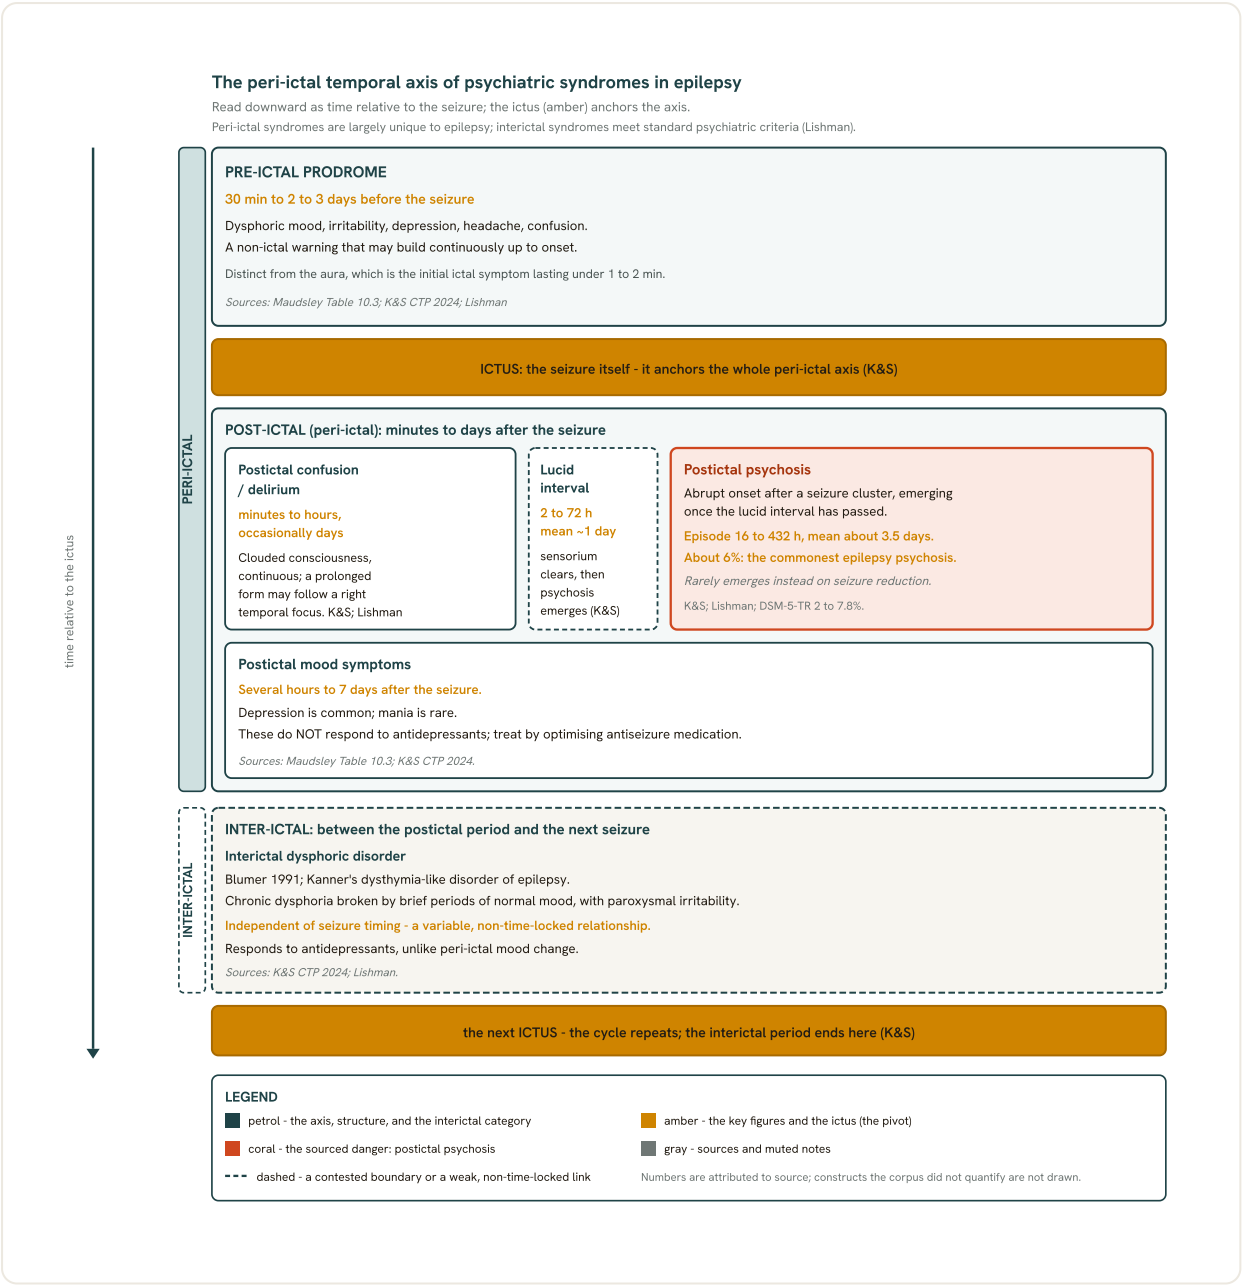


Fig 9.1 **The peri-ictal axis.** Psychiatric syndromes of epilepsy arranged by their timing relative to the seizure. Peri-ictal syndromes (the pre-ictal prodrome, ictal states, and postictal states, including postictal psychosis after a lucid interval) cluster around the ictus and are largely unique to epilepsy, whereas interictal syndromes such as the interictal dysphoric disorder are not time-locked and meet standard psychiatric criteria.

9.1 The vocabulary of timing: ictus, peri-ictal, interictal

Fix the reference points before anything is placed against them. The classification cannot be used until its temporal terms are pinned down, because each names a window in time rather than a symptom. The **ictus** is the seizure event itself; it is the zero point of the timeline, and Kaplan and Sadock's Comprehensive Textbook of Psychiatry treats it as the anchor from which the entire temporal axis is drawn. The **interictal period** is the stretch that lies between the settling of one seizure's aftermath and the onset of the next, which is to say the ordinary baseline state in which a person with epilepsy spends the great majority of life. The **peri-ictal period** is the narrow band just before or just after the ictus, and it carries a deliberate looseness in its very definition: it is the term reached for precisely when there is not enough information to say exactly where the seizure begins or ends.

That looseness is not sloppiness but clinical honesty. The electrical event has a beginning and an end that are frequently invisible without simultaneous recording, so a behavioural change that seems to precede a convulsion may in fact be the earliest, purely subjective fragment of the discharge, and an apparent aftermath may be the tail of ongoing seizure activity. Naming such a phenomenon peri-ictal, rather than committing it to ictal or postictal, is an open admission that without electrographic confirmation the exact boundary is unknown. This is the real work the peri-ictal label does: it holds together the phenomena that hug the seizure without forcing a precision the clinic rarely possesses. The interictal state, by contrast, is defined by its distance from the fit. It is what remains once the postictal disturbances have cleared and before the next attack announces itself, and it is against that quiet baseline that any enduring, seizure-independent psychiatric disorder must be judged.

9.2 The categories, by relationship to the seizure

The spine of the whole topic. With the vocabulary fixed, the psychiatric manifestations of epilepsy sort into **four** categories defined solely by their relationship to the ictus. Kaplan and Sadock's Comprehensive Textbook of Psychiatry sets them out as the **ictal** phenomena, which are the psychiatric expression of the discharge itself and of which the psychic auras are the paradigm; the peri-ictal phenomena clustered around the fit, of which postictal confusion is the everyday example; the interictal disorders arising in the baseline state, of which an interictal psychosis is representative; and a fourth, variably related group whose binding to the discharge is loose, and into which most mood disorders fall. The categories are exhaustive and they are mutually exclusive at the level of a single episode, which is exactly what makes the scheme usable at speed. The ictal category is reserved for a phenomenon that is unambiguously the discharge itself; where the boundary between the discharge and its aftermath cannot be drawn, the phenomenon is placed in the peri-ictal category rather than the ictal one, and it is that rule that keeps the two from overlapping.

CATEGORY	TIMING RELATIVE TO THE ICTUS	REPRESENTATIVE DISTURBANCE	DIAGNOSTIC FOOTING
Ictal	during the discharge	psychic auras	part of the seizure itself
Peri-ictal	the hours just around the fit	postictal confusion	largely unique to epilepsy
Interictal	the baseline between attacks	interictal psychosis	meets standard psychiatric criteria
Variably related	loose or uncertain tie to the fit	most mood disorders	relationship to the discharge not established

The categories of psychiatric disturbance in epilepsy, ordered by their temporal relationship to the seizure; each representative disturbance is developed as its own topic elsewhere in this chapter and is named here only to locate it.

The ordering is not arbitrary. Read from the discharge outward, each category is one step further removed from the electrical event, and that distance tracks a real shift in what the phenomenon means and what to do about it. An ictal disturbance is the seizure wearing a psychiatric face; a peri-ictal one is the brain's disturbed state in the immediate orbit of the fit; an interictal one is a

disorder living in the quiet between attacks; and the variably related group is the honest acknowledgement that some disturbances refuse to keep a schedule. Holding these in mind as a single ordered axis, rather than as unrelated headings, is what lets a clinician move from a raw presentation to a working category in a single step.

RULE **Classify by timing before you classify by symptom.** In epilepsy the same cross-sectional picture can arise in any of the categories, so the phenomenology alone will not place it. The relationship to the seizure is the primary sort key, and it must be established first, because it, and not the symptom list, decides both the diagnosis and the first treatment move.

MNEMONIC

AXIS

The positions on the axis, held as one word. **A**round the fit is the peri-ictal band. **X** marks the discharge itself, the ictal category. **I**n between the attacks is the interictal baseline. **S**lackly bound, keeping no reliable schedule, is the variably related group where most mood disorders sit. Each letter asserts only a timing this topic states.

9.3 What the peri-ictal category gathers together

One category, several kinds of phenomena. The peri-ictal bin is the busiest of the categories and the one most often misread, because it is not a single syndrome but a gathering of **three** temporally distinct groups. Kaplan and Sadock's Comprehensive Textbook of Psychiatry, in its table of behavioural disorders in epilepsy, collects within the peri-ictal category the prodromal disturbances that precede the fit, the postictal disturbances that follow it, and the boundary-blurred phenomena in which it cannot be said cleanly whether a disturbance is still the discharge or already its aftermath.

- The **prodromal** phenomena, such as the dysphoria and irritability that can build over hours to a few days before a seizure, taken up as the preictal prodrome in its own dedicated discussion.
- The postictal phenomena that emerge as the discharge subsides, spanning the confusion and delirium of the immediate aftermath through to the postictal psychoses that declare themselves only after a symptom-free interval.
- The boundary-blurred phenomena, including the peri-ictal psychoses, where the very lack of information about when the seizure starts and stops is what forces the peri-ictal label onto the picture.

Grouping these under one heading is not a filing convenience; it reflects a shared biology. Each is a direct product of a brain destabilised by the seizure or by the state immediately surrounding it, which is why they arrive stereotyped, run a limited course, and settle as that state resolves. It also explains why the peri-ictal category behaves so differently from the interictal one at the level of workup and treatment, a difference taken up next. Keeping the prodromal, postictal and mixed subgroups mentally separate, while recognising that they share the peri-ictal home, prevents the

common error of treating everything that happens near a seizure as a single undifferentiated event.

9.4 Peri-ictal versus interictal: unique to epilepsy, or not

The split that carries most of the diagnostic weight. The sharpest and most examinable line on the axis runs between the peri-ictal and the interictal categories, and Lishman's Organic Psychiatry frames it in a way worth memorising. The peri-ictal presentations, the prodrome and the postictal psychosis among them, are largely unique to epilepsy: they have no clean counterpart in general psychiatry, so their very form points back toward the seizure and is close to diagnostic of an epilepsy-related state. The interictal disorders, by contrast, meet the standard psychiatric criteria used in patients who have no epilepsy at all, and they look like the disorders seen everywhere else. That contrast has a direct clinical consequence. A peri-ictal presentation is essentially telling you its own origin, whereas an interictal presentation must be worked up as you would work up any psychiatric disorder, with the epilepsy standing as context and vulnerability rather than as the proximate cause. The two categories demand different reasoning even when their surface phenomenology briefly resembles, and collapsing them loses the point of the axis.

TRAP **Treating every psychiatric disorder in epilepsy as epilepsy-specific.** Candidates lose the mark by assuming that because a disturbance occurs in a patient with epilepsy it must be a peri-ictal, seizure-bound phenomenon. The interictal disorders meet ordinary psychiatric criteria and are diagnosed exactly as they would be in anyone else. Reserve the epilepsy-unique framing for the peri-ictal band, and keep the interictal disorders on standard footing.

9.5 The fourth category: disturbances only loosely tied to the fit

Not everything dates cleanly to the seizure, and the scheme says so. The fourth category is easy to forget precisely because it is defined by the absence of a firm temporal tie, yet naming it is itself examinable. Kaplan and Sadock's Comprehensive Textbook of Psychiatry places into this **variably related** group the disturbances whose relationship to the discharge is weak or not established, and it notes that most mood disorders belong here. A depressive or anxious presentation in epilepsy sometimes clusters before or after seizures, sometimes runs steadily in the interictal baseline, and often shows no dependable relationship to the fits at all, which is exactly why the classification declines to force it into a single slot. The honesty of this fourth bin is its strength. It resists the reflex, common in candidates and clinicians alike, to assume that every disturbance in epilepsy must lock onto the seizure, and it leaves room for the mood syndromes that genuinely do keep a schedule, such as the postictal mood changes and the interictal dysphoric disorder of epilepsy, to be recognised and taught in their own right elsewhere in this chapter. What sits in the variable category is the broad mass of mood disturbance that will not be pinned to the discharge, and treating that mass as though it were peri-ictal is a category error the axis is designed to prevent.

9.6 Peri-ictal psychosis and the paradox of seizure frequency

A disturbance that can worsen when seizures get either worse or better. The peri-ictal psychoses illustrate a counter-intuitive feature that the axis makes sense of. As a rule, peri-ictal disturbance tracks seizure burden, so that a run of fits is the usual precipitant and better seizure

control brings better behaviour. But the **peri-ictal psychosis** can move in either direction: it may worsen as seizure activity increases, or, less commonly, it may emerge precisely as seizures come under control and the record quietens. That second, decreased-frequency form is the bridge to forced normalisation, the phenomenon in which behavioural disturbance appears as seizure activity is suppressed, which is developed in its own dedicated discussion elsewhere in this chapter. The clinical lesson is that a psychosis appearing on the axis cannot be read off seizure frequency alone; the direction of change matters as much as the fact of it, and improving the electroencephalogram is not always the same as improving the patient.

PITFALL

When a patient's seizures have just been brought under good control and psychosis then appears, the reflex is to read it as an unrelated new illness and to press on with the successful anticonvulsant. In the decreased-frequency form the improving seizure control may itself be driving the disturbance, so intensifying suppression can make matters worse. A psychosis that arrives as the fits recede should prompt a rethink of the regimen, not an automatic escalation of it.

9.7 The management principle the timing dictates

The axis is not only diagnostic; it sets the first move in treatment. Once a psychiatric disturbance in epilepsy has been placed in the peri-ictal band, its management follows a principle that the timing itself dictates, and the guidelines that govern it are, for the seizures, the International League Against Epilepsy and the National Institute for Health and Care Excellence epilepsy guidance, and, for the psychiatric layer, the epilepsy-specific psychiatric-management literature together with the Maudsley Prescribing Guidelines. The organising claim is blunt and high-yield: every peri-ictal psychiatric symptom is first treated by optimising the **antiseizure medication**, whether the symptom is preictal, postictal or of blurred boundary, because there is no evidence that psychotropic drugs prevent the peri-ictal symptoms themselves.

The **LOCUS** of that care is usually the outpatient clinic, where the antiepileptic regimen is reviewed and tightened; it escalates to the inpatient ward or the emergency department only when a peri-ictal state such as a florid postictal psychosis or severe agitation threatens safety. The **focus** shifts by phase: the acute target is the disturbed behaviour of the current episode and the risk it carries, the short-to-mid-term target is the seizures that generated it, and the long-term aim is durable seizure control so that the peri-ictal disturbances do not recur. The **modus** is led by the pharmacological mode, but pointed at the seizures rather than at the symptom, optimising the anticonvulsant before reaching for any psychotropic; the psychological, social and procedural modes support this work, through education, safety planning and attention to the precipitants of seizure clusters, but none of them substitutes for getting the seizure treatment right. The whole plan is an application of the axis: because the disturbance is peri-ictal, the lever that moves it is the one that moves the seizures.

IN INDIA

The rule that peri-ictal symptoms are managed first by optimising seizure control lands with particular force in Indian practice, where the epilepsy treatment gap is wide and many patients reach the psychiatrist with poorly controlled or unmonitored seizures. Inexpensive agents such as phenobarbital remain in first-line use in resource-limited settings, and video-electroencephalography to confirm the exact timing of an episode is available only in tertiary centres, so the peri-ictal placement is usually a clinical judgement made from a careful history. In this setting, getting the anticonvulsant regimen right is frequently the single most effective psychiatric intervention available, and it is reachable without the specialist infrastructure that the psychotropic-first reflex would otherwise wait for.

9.8 Putting the axis to work at the bedside

One question, asked first, reorganises the whole assessment. In practice the axis is used as a first filter. Faced with any psychiatric presentation in a patient with epilepsy, the clinician asks where it sits in relation to the last seizure and the next, and the answer routes the entire workup. A phenomenon locked to the discharge points toward the ictal syndromes, ictal fear and the ictal affective phenomena, ictal psychosis and the ictal depressive phenomena. A phenomenon hugging the fit points into the peri-ictal band, the preictal prodrome of dysphoria and irritability, postictal confusion and delirium, postictal psychosis with its lucid interval, and postictal depression, mania and anxiety. A phenomenon settled in the quiet baseline points to the interictal disorders, the interictal dysphoric disorder of epilepsy and the schizophrenia-like psychosis of epilepsy. And a mood disturbance that will not date cleanly sits, honestly, in the variable category. The accompanying figure lays the same axis out visually for revision.

Held this way, the axis earns its marks by doing several things in one move:

- It narrows a wide differential to a short list of candidate syndromes, because each position on the axis carries only a handful.
- It separates the epilepsy-unique peri-ictal states from the interictal disorders that meet ordinary psychiatric criteria, which is the split that most often decides the diagnosis.
- It fixes the first treatment move, since a peri-ictal placement points straight at optimising seizure control rather than at a psychotropic.

GOOD TO REMEMBER

Before reaching for an antidepressant, antipsychotic or anxiolytic in a patient with epilepsy, timestamp the symptom against the seizure. Where it sits, coincident with the discharge, hugging the fit, in the quiet baseline or floating free, changes the diagnosis you reach and the very first prescription you write. The timing is not background detail; it is the sort key for the whole assessment.

Ictal	Peri-ictal	Interictal	Variable
THE DISCHARGE ITSELF	AROUND THE FIT: PRODROME, POSTICTAL	THE BASELINE BETWEEN ATTACKS	MOST MOOD DISORDERS, NO FIXED TIE

Every psychiatric disturbance in epilepsy is classified first by its timing to the seizure, ictal, peri-ictal, interictal or only variably related, and that timing, not the symptom in cross-section, is what sets both the diagnosis and the first move in management.

10 · Ictal fear, ictal panic and ictal affective phenomena

The fear that opens a seizure is an examiner's favourite, and with reason. Among the peri-ictal psychiatric phenomena, the syndromes sorted by their timing relative to the discharge, ictal affect is the material that arises during the electrical event itself, and of all ictal affects the abrupt surge of dread is both the commonest and the one most readily mistaken for a primary psychiatric disorder. A candidate is expected to recognise it on sight, localise it, and lift it out from the interictal anxiety disorders that also afflict people with epilepsy. This topic takes the ictal slot on the temporal classification of peri-ictal syndromes and fills it with the affective phenomenology that lives there; it names that classifying axis but does not re-derive it, because the framework is developed as its own topic. What follows is the phenomenology of ictal fear and panic, the wider palette of ictal affect around them, and the single differential that carries most of the marks.

10.1 Ictal fear as a psychic aura

Ictal fear is fear generated by the seizure discharge itself, experienced as a **psychic aura** rather than as a reaction to anything in the patient's surroundings. Its intensity is not fixed. The same underlying mechanism yields a faint, barely nameable unease in one patient and overwhelming terror in another, so the phenomenon is best held in mind as a spectrum that runs from vague apprehension at one pole to abject fright at the other, as it is set out in Kaplan and Sadock's Comprehensive Textbook of Psychiatry. The clinically decisive fact is that this fear can constitute the whole seizure. It may occur with no motor sign, no automatism and no march into anything more obviously epileptic, which is precisely why it so often reaches a psychiatrist as an unexplained anxiety symptom long before epilepsy is entertained. The subjective report is frequently hard to put into words, and a patient may reach for figurative language, a sense of impending doom, of something terrible about to happen, of the ground giving way, none of which maps onto a real event; it is this ineffable, causeless quality that both defines the experience and marks it as arising from within rather than from the world.

Fear of this ictal type is also the commonest **affective aura** of temporal lobe epilepsy, and its defining temporal signature, described in Lishman's Organic Psychiatry, is that it wells up suddenly and without provocation. A patient who is calm and unstressed is abruptly gripped by dread that has no object and no trigger, and then it recedes of its own accord. That combination, maximal at onset, unprovoked, and self-limiting, is the phenomenological fingerprint that will later separate an ictal panic from a psychiatric one, and it repays being fixed firmly at the outset. Temporal lobe epilepsy supplies the substrate for this and for much of the affective phenomenology of seizures, but its anatomy and semiology are developed as their own topic; here it is enough to know that the fear and the temporal lobe travel together, and that a psychic aura of pure dread is a temporal-lobe event until something else is shown.

10.2 The ictal amygdala phenomenon and the fear network

The reason ictal fear binds so tightly to the temporal lobe is anatomical. The affect is manufactured by discharge in the medial temporal structures, and above all in the amygdala, the limbic node whose physiological task is the tagging of salience and threat; when a seizure recruits it, the output is fear with no frightening world to justify it. Lishman's Organic Psychiatry names this the **ictal amygdala phenomenon**, and it is the mechanistic heart of the topic. The wider structural account of the limbic system, its nuclei and its circuits, belongs to the Functional Neuroanatomy and Neurophysiology notes and is not rehearsed here; the point that matters clinically is that the fear is produced by the brain, not perceived by it, so there is no proportionate threat to be found and none should be sought. This also accounts for the sheer conviction the affect carries. Because it is generated by the same machinery that ordinarily signals real danger, the patient does not experience it as an idea about fear but as fear itself, fully embodied, which is part of why it is so distressing to endure and so persuasive to the clinician who has not yet placed it.

RULE **Sudden unprovoked fear that is the entire seizure points to the medial temporal lobe.** The commonest affective aura of temporal lobe epilepsy arises from amygdala discharge, so an isolated psychic aura of dread is, until shown otherwise, a mesial temporal event and not a mood, a nerve or a reaction to circumstance. Localisation follows from the phenomenology before any recording is made.

Intense fear, however, is not purely amygdalar. As the seizure grows, the affect maps onto a **distributed limbic network** that draws in the medial temporal lobe together with orbitofrontal cortex and the anterior cingulate, and the clinical escalation from a mild change in mood to frank terror tracks the spread of the discharge on the electroencephalogram towards the frontal regions. This yields a useful teaching image: the strength of the fear is, in effect, a readout of how much cortex the seizure has captured. It also explains why the most severe ictal fear tends to accompany the seizures that recruit the widest network rather than the briefest, most contained auras, and why a patient's worst episodes may look and feel categorically different from the mild apprehension that heralds a milder attack, even though both are the same phenomenon at different amplitudes.

10.3 Behavioural accompaniments and the paradox of preserved awareness

What the observer sees is as diagnostic as what the patient feels. When ictal fear reaches high intensity it drives a small, recognisable set of behaviours that Lishman's Organic Psychiatry sets out, and knowing them lets a clinician read the episode from the outside even when the patient cannot report it in the moment. The accompaniments cluster into three forms:

- a sudden call for help, the patient turning to or calling out to whoever is present;
- marked agitation, restless and driven rather than purposeful;
- **frightened immobility**, the patient frozen and mute, held still by terror.

The feature that catches candidates out is that through all of this awareness is usually preserved. The patient is frightened, may cry out, looks about, and can afterwards describe the experience in

detail, all while the seizure is in progress. This is the expected state of a focal event that has not yet impaired consciousness, and it is the source of a recurring clinical error, because the retained, reportable quality of the episode is read as evidence against a seizure rather than as a normal property of one.

PITFALL

Because the patient can call out, scan the room and later narrate the whole episode, clinicians conclude that it "cannot be a seizure" and reach instead for a functional or a purely anxious explanation. Preserved awareness is the expected accompaniment of this kind of focal event, not an argument against it, and treating retained consciousness as exclusionary is a well-worn route to a missed diagnosis.

10.4 Beyond fear: the wider menu of ictal affect

Fear is the loudest ictal affect but it is not the only one, and a complete account of the topic has to name the rest of the palette. At the opposite emotional pole are the pleasurable auras, transports of joy, elation and even ecstasy; these are the classical **ecstatic auras** associated with Dostoyevsky, described in Kaplan and Sadock's Comprehensive Textbook of Psychiatry, and while they are genuine seizure phenomena they are distinctly less frequent than fear. Their importance out of proportion to their rarity lies in their capacity to mislead: a patient who describes a recurring rapture, a sense of revelation or of perfect wellbeing, may never think to mention it as a symptom at all, and the clinician who does not ask will not hear of it. In the older literature such states drew attention out of proportion to their frequency because of their link to religious and mystical experience, and the practical residue of that history is a standing reminder to take an account of transcendent episodes seriously as a possible seizure rather than dismissing it as fancy or as devotion.

The unpleasant end of the palette also extends past fear. Lishman's Organic Psychiatry records that a seizure can manufacture depression and **ictal guilt**, and, rarely, anger, so the affective aura is capable of counterfeiting a fair range of primary mood states in miniature.

- pleasurable auras of joy, elation or ecstasy, less common than the fearful ones;
- ictal depression and guilt, momentary and tied to the discharge;
- ictal anger, which is rare.

Two boundaries need holding here. The depressive end of this menu, the ictal depressive phenomena proper, is developed in full alongside the ictal psychoses as its own topic and is only named here, not taught. And these ictal affects are all momentary, locked to the electrical event, and must not be confused with the sustained interictal dysphoric disorder of epilepsy, a separate, enduring mood syndrome that keeps its own company elsewhere in this chapter. The unifying feature of everything in this section is transience: an ictal affect, whether dread, ecstasy or guilt, is the seizure wearing an emotion, and it lifts as the discharge passes.

10.5 Where ictal affect sits among ictal psychiatric symptoms

The ranking is worth memorising because it predicts what the clinic actually delivers. Standard tabulations of ictal psychiatric symptomatology, such as the one in the Maudsley, arrange the phenomena by frequency, and the order is stable and high-yield. Ictal fear and **ictal panic** head the list as the most common presentation, ictal depressive symptoms come next, and ictal psychosis is rare. Read as a single hierarchy, this tells the psychiatrist that when a seizure declares itself through a psychiatric symptom, the odds overwhelmingly favour fear or panic, and it is the presentation to expect and to be ready to recognise.

Commonest	Next	Rare
ICTAL AFFECT: FEAR OR PANIC	ICTAL DEPRESSIVE SYMPTOMS	ICTAL PSYCHOSIS

Both the depressive and the psychotic ends of that ranking belong to the neighbouring topic on ictal psychosis and ictal depressive phenomena, and neither is developed here. What matters for the present topic is that fear and panic sit unambiguously at the top of the list, so that recognising an ictal affect is, most of the time, the work of recognising ictal fear. The hierarchy also carries a quiet warning against the reverse inference: the rarity of ictal psychosis does not make it impossible, and a strange, brief, stereotyped psychotic fragment in a person with epilepsy deserves the same timing question that fear does. For the trainee the corollary is a habit rather than a fact: when epilepsy and a psychiatric symptom coincide, expect fear first, and let that expectation shape the very first questions asked about the timing of the episode, its stereotypy and the company it keeps.

10.6 Ictal panic versus a primary panic disorder

This is where the marks are won and lost. The single most examined point in this topic is the separation of an ictal panic from a primary, interictal panic or anxiety disorder, because the two can feel identical to the patient and yet demand completely different management. The interictal anxiety disorders that people with epilepsy also develop are a genuine and separate problem, taken up among the anxiety disorders in epilepsy, and an ictal panic must be lifted out from among them before either is treated. The features that betray an ictal origin are its brevity, its stereotypy and its **epileptiform semiology**: an ictal panic is short and self-limiting, replays in almost the same form each time, and carries the hallmarks of a seizure, so that it may cloud consciousness or evolve into recognisable seizure activity, whereas a psychiatric panic attack runs longer, varies from episode to episode, and leaves consciousness entirely clear. The stakes of getting this wrong run in both directions. A missed ictal panic leaves a focal epilepsy untreated and exposes the patient to the hazards that unrecognised seizures carry, while an ictal panic wrongly assumed in someone who in fact has an ordinary anxiety disorder invites an unnecessary anticonvulsant and delays the psychological treatment that would actually have worked.

FEATURE	POINTS TO AN ICTAL PANIC	POINTS TO A PRIMARY ANXIETY DISORDER
Onset	abrupt, unprovoked, maximal at the start	may build, often situational or cued
Course	brief and self-limiting	typically longer and more variable
Stereotypy	highly stereotyped, recurs near-identically	content and setting vary
Consciousness	may be impaired; automatisms may follow	clear throughout
Correlate	epileptiform semiology; other seizure signs may appear	no epileptiform correlate

Features that separate an ictal panic from a primary interictal anxiety disorder; the interictal disorders themselves are developed as their own topic and are named here only for the contrast.

TRAP **Labelling a brief, stereotyped ictal panic as panic disorder and treating it psychiatrically.** Candidates anchor on the shared subjective experience and prescribe an antidepressant or refer for cognitive therapy, leaving the underlying focal seizures untreated. The stereotypy and the epileptiform features, not the feeling, decide the diagnosis, and the feeling is the one thing the two conditions genuinely share.

MNEMONIC

Four S's: Sudden, Short, Stereotyped, Semiology

When a surge of fear is **S**udden and unprovoked, **S**hort and self-limiting, **S**tereotyped from one episode to the next, and carries epileptiform **S**emiology, it is a seizure until proven otherwise. Each letter asserts only a property this topic states of ictal fear and panic.

GOOD TO REMEMBER

Any patient whose "panic attacks" are brief, arrive out of nowhere, replay in nearly the same form every time, and are followed by a patch of amnesia deserves an electroencephalogram and a neurology opinion before the label of panic disorder is allowed to settle. The cost of asking is small; the cost of missing a focal epilepsy dressed as anxiety is not.

A brief, stereotyped, unprovoked surge of fear that recurs identically, and may cloud awareness, is a seizure until proven otherwise.

10.7 Recognition, and what follows from it

Because ictal fear and panic are seizures, their management is the management of the epilepsy, and the psychiatrist's decisive contribution is diagnostic: to recognise the phenomenon, name it correctly, and route the patient appropriately rather than to reach for a psychotropic. The guideline anchors for that work are the International League Against Epilepsy and the National Institute for Health and Care Excellence epilepsy guidance, supported by the epilepsy-specific

psychiatric-management literature; none of them treats ictal fear as a mood or an anxiety disorder in its own right, and all of them place the target on the seizure.

The **LOCUS** of care is overwhelmingly the outpatient clinic, shared between psychiatry and neurology or a dedicated epilepsy service; the emergency department enters only when an aura evolves into a prolonged or repetitive seizure, or when the terror drives behaviour that threatens safety. The **focus** shifts by phase: the first target is the diagnosis itself, and thereafter the symptom that must be controlled is the seizure, so that the affect settles as seizure control improves, with immediate safety the aim in the acute event, seizure reduction over the short and mid term, and durable prevention of recurrence over the long term. The **modus** runs across every available mode. The drug mode is antiseizure medication, whose selection belongs with the antiepileptic drugs used in epilepsy and is not rehearsed here; the psychological mode is psychoeducation and reassurance that reframes the experience as a seizure rather than as madness or a heart attack, which is often profoundly relieving in itself; the procedural mode, in drug-resistant focal epilepsy, is evaluation for epilepsy surgery; and the social mode covers the driving, employment and family counselling that any focal epilepsy demands. The organising idea is that the lever which moves the fear is the one that moves the seizures.

IN INDIA

A fearful or ecstatic psychic aura is readily read through a spiritual or possession framework, and families frequently consult faith healers or temples before a physician, so the seizure origin surfaces late. Limited and unevenly distributed access to electroencephalography compounds the delay, and stigma around both epilepsy and mental illness discourages disclosure of experiences that a patient may already fear will be taken as madness. The practical lesson is to ask directly, in anyone presenting with sudden unexplained terror or with recurring rapturous experiences, whether the episodes are brief and stereotyped, and to lower the threshold for referral rather than assume a primary psychiatric or spiritual cause.

11 · Ictal psychosis and ictal depressive phenomena

When the disturbance is the seizure, the treatment is to stop the seizure. Two of the loudest psychiatric presentations in epilepsy, a florid psychosis and a profound depression, can each occur as the direct expression of an ongoing electrical discharge. When they do, they sit in the ictal slot of the peri-ictal axis, reserved for phenomena that are not merely provoked by a seizure but are the seizure, worn as a psychiatric face. This is what makes **ictal psychosis** and **ictal depression** distinct from every syndrome that follows a fit across a gap. The marks are here because the ictal versions are the most dangerous to misread: a psychosis that is actually a status is worsened by the antipsychotic reflex, and a depression that is actually a discharge can drive a patient to self-harm while the clinician waits for a mood disorder to declare itself.

These presentations belong to the temporal classification of peri-ictal psychiatric syndromes, developed as its own framework elsewhere in this chapter, and their defining feature is the one the axis makes visible: there is no interval between the discharge and the disturbance because they are the same event, which dictates a management logic closer to neurology than to

psychiatry. The account below teaches what an ictal psychosis and an ictal depression actually are, why each is missed, and the load-bearing distinction that keeps them apart from their postictal mimics. The ictal affective phenomena proper, ictal fear and ictal panic, are taken up in their own topic and are only named here.

11.1 The raw material: ictal psychic phenomena

A focal discharge can generate experience, not just movement. A focal seizure arising in or spreading through limbic and temporal cortex need not announce itself as a convulsion; it can present entirely as altered inner experience. Kaplan and Sadock's Comprehensive Textbook of Psychiatry gathers these under the heading of **ictal psychic phenomena**, the substrate from which both an ictal psychosis and an ictal depression are built.

- Mood changes, in which the affect itself is generated by the discharge rather than reflecting the patient's circumstances.
- Derealization and depersonalization, the sense that the surroundings or the self have become unreal or detached.
- **Forced thinking**, the intrusion of a stereotyped, unbidden thought the patient cannot resist and does not recognise as their own initiative.

What unites these is their origin. Each is a fragment of experience manufactured by abnormal cortical activity, which is why they arrive suddenly, feel alien, and are stereotyped from one seizure to the next. When the discharge is brief, these phenomena are the psychic auras a focal seizure can carry; when it is sustained, the same machinery produces something that looks less like a fleeting aura than like a psychiatric illness in its own right. That transition, from a momentary psychic symptom to a sustained state driven by continuous seizure activity, is exactly what an ictal psychosis is.

11.2 Ictal psychosis is nonconvulsive status epilepticus

The identity that unlocks the whole topic. The single most important thing to grasp is that ictal psychosis is not a psychosis that happens to accompany a seizure; it is a seizure that happens to look like a psychosis. Kaufman's Clinical Neurology for Psychiatrists makes the point directly: the behaviour of a **nonconvulsive status epilepticus** can be so bizarre, disorganised and psychiatric in flavour that it earns the label of ictal psychosis. There is no separate disease to diagnose. The patient is in continuous or near-continuous seizure activity that never crosses into visible convulsion, and the disturbance the clinician observes is the outward face of that ongoing discharge.

Framing it this way reorganises everything downstream. If the psychosis is the discharge, then it is concurrent with the discharge, persists for as long as the discharge persists, and resolves the moment the discharge is aborted, none of which is true of a psychiatric illness that merely coexists with epilepsy. It also explains the strangeness of the presentation: depending on where the activity sits and how it spreads, nonconvulsive status can produce anything from a fluctuating confusional state to a densely psychotic picture with disorganised behaviour. Holding ictal psychosis as an electrographic event rather than a mental illness is the conceptual move that

makes diagnosis, timing and treatment fall into place, and the one candidates most often fail to make.

11.3 How long it lasts and how it looks

Sustained, fluctuating, and cognitively disorganised. Because the disturbance is a continuing discharge rather than a passing aura, it is not measured in seconds. Kaufman's Clinical Neurology for Psychiatrists describes the aberrations of nonconvulsive status running across **hours to days**, a duration that marks the state out from an ordinary focal seizure and gives it the sustained quality that makes it resemble a primary psychiatric disorder. The picture is dominated by disorganisation. The same source characterises the state by thought disorder, impairment of language, and an **altered sensorium**. The patient is therefore muddled, disconnected and hard to reach rather than cleanly psychotic in the way a person with schizophrenia is, with awareness that waxes and wanes.

That combination is diagnostically useful. A florid but clear-sensorium psychosis in a fully alert patient sits uneasily with ictal psychosis; a fluctuating, disorganised state in which language and coherence come and go, in a person with known epilepsy, sits comfortably with it. The waxing and waning itself is a clue, because the mental state tracks the shifting intensity of the discharge. That texture is quite different from the loud, affectively coloured, clear-headed psychosis that arrives after a lucid interval in the postictal form.

MNEMONIC

BEND

The identity of ictal psychosis in four beats. **B**enzodiazepine given intravenously is the first move, aborting the status and lifting the mental state. **E**EG makes the diagnosis, because the psychosis is a discharge and nothing else. **N**o lucid interval separates the seizure from the disturbance; they are one event. **D**ischarge itself is the disorder, a nonconvulsive status wearing a psychiatric face. Each letter asserts only what this topic states.

11.4 Why it is missed: the mimics

Everything it resembles is a psychiatric or neurological emergency of its own. The danger of ictal psychosis lies less in the disorder itself, which is eminently treatable once recognised, than in the company it keeps. Kaufman's Clinical Neurology for Psychiatrists notes that it readily masquerades as several other conditions, each of which pulls the clinician toward a different and often wrong response.

- It looks like delirium, a clouded, fluctuating, disorganised state, and the reflex is to hunt for a metabolic or infective cause while the discharge goes unrecorded. Delirium in its own right, its assessment and subtypes, belongs to the Stroke, TBI, delirium and focal brain syndromes notes.
- It looks like the aftermath of traumatic brain injury, especially where a seizure has caused a fall or head injury and epilepsy coexist, so the disturbance is attributed to structural damage rather than a live electrical event.

- It looks like an acute psychotic episode, and here the error is most costly, because the instinct is to reach for an antipsychotic and admit the patient as a first psychiatric break.

The common thread is that every mimic invites a plan built for a different disease, and none includes the one investigation that settles the matter, so a delirium workup, a brain-injury attribution or a first-episode-psychosis admission can proceed for hours while the patient remains in status. The lesson is not that ictal psychosis is hard to treat but that it is easy to overlook, and the cost of overlooking it is prolonged, unrecognised seizure activity. Where a psychosis in a person with epilepsy is disorganised, fluctuating and clouded, the possibility that it is a discharge should be actively excluded.

TRAP **Treating the status as a functional psychosis.** The classic error is to see a disorganised, agitated, psychotic patient with epilepsy and start an antipsychotic, without an electroencephalogram, as though this were an acute psychiatric break. That misses the diagnosis, delays the only effective treatment, and reaches for a drug class chosen for the wrong disease. The corrective is a reflex: disorganised psychosis plus epilepsy means record the brain before you sedate the mind.

11.5 The load-bearing distinction: no lucid interval

The one line that separates ictal from postictal psychosis. If a single fact from this topic is worth carrying into an examination, it is the temporal relationship each psychosis holds to the seizure. Kaufman's Clinical Neurology for Psychiatrists and Kaplan and Sadock's Comprehensive Textbook of Psychiatry converge on the discriminator: the ictal psychosis is concurrent with the ongoing discharge and has no lucid interval at all, whereas the postictal form declares itself only after a symptom-free gap of apparent recovery. The **lucid interval** is the pivot on which the two diagnoses turn. In the ictal form the patient is disturbed continuously from within the seizure, with no window of return to clarity, because the process has never stopped. In the **postictal psychosis**, developed in full as its own topic, the seizure ends, the patient appears to recover, and only then does the psychosis erupt.

FEATURE	ICTAL PSYCHOSIS	POSTICTAL PSYCHOSIS
Relation to discharge	concurrent with ongoing discharge	after the seizure has ended
Lucid interval	**none**	present, a symptom-free gap
Underlying process	nonconvulsive status epilepticus	a seizure-driven aftermath state
Diagnosis	electroencephalogram shows the discharge	reconstructed from the timing
First-line treatment	intravenous benzodiazepine to abort the status	self-limiting; antipsychotic or benzodiazepine as needed

The ictal versus postictal psychosis distinction, the load-bearing separation in this part of the chapter; the postictal entity, its lucid interval and its numbers are taught in full in its own topic and are named here only to draw the contrast.

The distinction matters because it changes the first thing the clinician does. Read the psychosis as ictal and the priority is to record the brain and abort a status; read it as postictal and the priority

is safety and containment of a self-limiting episode. Getting the timing wrong sends the patient down the wrong pathway from the outset, which is why the discriminator should be stated before any description of symptoms.

RULE **Ictal psychosis is the discharge; postictal psychosis follows it.** The ictal form coincides with ongoing seizure activity and has no lucid interval, so it is diagnosed on the electroencephalogram and terminated by aborting the status; the postictal form arrives after a gap of apparent recovery and is a separate syndrome. Every decision in this topic descends from that one structural difference.

11.6 Ictal depression and its danger

A discharge can be low, not just frightened or elated. The affective side of the ictal slot is easy to underplay, because a brief ictal mood change can seem trivial next to a psychosis. It is not, for two reasons Kaplan and Sadock's Comprehensive Textbook of Psychiatry sets out. First, unlike most ictal phenomena that end when the seizure ends, an ictal depression can extend for **days or longer** after the seizure has passed, so the low mood outlives the discharge that generated it and comes to look like an ordinary depressive episode. Second, and this carries the weight, though it is rare an ictal depression can outlast the ictus and, at its most dangerous, drive the patient to suicide. Such a state is a genuine risk, because its severity can be out of proportion to any situational trigger and it can persist long enough to act on.

The significance is that an ictal depression breaks the comfortable assumption that ictal phenomena are always brief and self-cancelling. Most are; this one need not be. A patient whose depression is the tail of a seizure discharge may present days later with a mood state that is severe, endogenous in quality and dangerous, and the epileptic origin is easy to overlook once the seizure is out of the immediate history. Depression in epilepsy and its screening, and the general question of suicide risk in epilepsy, are each developed elsewhere in this chapter; what is owned here is that a depressive state can be an ictal event in itself and carry lethal risk while it lasts.

PITFALL

A severe low mood in the days after a seizure is easily filed as an understandable reaction to having fits, or a mild postictal dip, and its danger discounted on the grounds that seizure-related mood changes are brief. An ictal depression can persist well beyond the fit and carry a real risk of suicide. Do not let the epileptic context soften the risk assessment; a persistent, severe depressive state after a seizure deserves the same safety scrutiny as any other, and arguably more.

11.7 Management

Here the psychiatrist thinks like a neurologist. The management of ictal psychosis is unusual because the target is not the mental state but the seizure that produces it, and the governing guidance is drawn accordingly from the International League Against Epilepsy and the National Institute for Health and Care Excellence epilepsy guidelines for the seizure, together with the epilepsy-specific psychiatric-management literature for the behavioural layer. The organising claim from Kaufman's Clinical Neurology for Psychiatrists is that nonconvulsive status is diagnosed on the electroencephalogram and treated first with an intravenous benzodiazepine, a

single move that both aborts the status and awakens the patient, so diagnosis and treatment are demonstrated in the same act.

The **LOCUS** of care is not the outpatient clinic. Because ictal psychosis is a nonconvulsive status, it is a neurological emergency requiring the emergency department or an inpatient setting with electroencephalography, intravenous access and monitoring; a disorganised, clouded patient cannot be safely held in the community while a status runs unrecorded. The **focus** shifts by phase: the acute target is to confirm and abort the discharge, the short-to-mid-term target is to find what precipitated the status and prevent recurrence by optimising the antiseizure medication, and the long-term aim is durable seizure control. The **modus** is led by the pharmacological mode but aimed at the seizure: the intravenous benzodiazepine first, then optimisation of the antiepileptic regimen, whose seizure-type-led selection is a topic in its own right and is not reopened here; the procedural mode, the electroencephalogram, is both the diagnostic instrument and the measure of response. This is the sharpest contrast with the postictal psychosis, which is self-limiting and contained with an antipsychotic or benzodiazepine: in the ictal form an antipsychotic is not the answer, and the question of which antipsychotic is safe against the seizure threshold belongs to its own topic. For an ictal depression the same principle applies once the discharge is settled, but the psychological and social modes come forward, because the acute priority is protection from the suicide risk the state can carry.

IN INDIA

The diagnosis rests on capturing the discharge, yet access to electroencephalography, let alone urgent or continuous recording, is concentrated in tertiary centres. In much of Indian practice such a patient reaches a clinician who cannot immediately record the brain, and the risk is that a nonconvulsive status is labelled an acute psychosis and started on an antipsychotic. The practical safeguard is a high index of suspicion and a low threshold for a therapeutic trial of an intravenous benzodiazepine, which is widely available and inexpensive, when the psychotic state is clouded rather than clear; a rapid awakening is both treatment and diagnosis. Where the treatment gap leaves seizures poorly controlled, the sustained discharges underlying ictal states are correspondingly more common.

11.8 Placing both at the bedside

One question routes both syndromes. A clouded psychosis and a persistent depression look nothing alike, yet they are managed by the same first question: is this the seizure, or something the seizure left behind. For the psychosis, concurrent-with-the-discharge means a status to abort, and the absence of any lucid interval confirms it. For the depression, an ictal origin means a mood state that can outlast the fit and turn lethal, deserving active safety measures rather than reassurance. Both answers come from timing the disturbance against the discharge, the discipline the temporal classification of peri-ictal syndromes teaches for the whole family. The differential is short: postictal confusion and the postictal psychosis, which follow the seizure across a gap; and, for the mood side, the interictal dysphoric disorder of epilepsy and the general depression of epilepsy, which are not ictal at all. The accompanying figure locates both ictal syndromes on the peri-ictal axis for revision.

Hours to days	None	Days or longer	IV benzodiazepine
DURATION OF ICTAL PSYCHOSIS	LUCID INTERVAL IN THE ICTAL FORM	AN ICTAL DEPRESSION CAN PERSIST	FIRST-LINE, ABORTS THE STATUS

GOOD TO REMEMBER

Any prolonged, fluctuating, clouded psychotic state in a person with epilepsy should prompt an electroencephalogram to exclude a nonconvulsive status before a psychiatric diagnosis is committed to, because the treatment is to abort a seizure, not to start an antipsychotic. And a severe depression that appears or persists in the days after a fit should be taken as seriously as any other, because it can be ictal in origin and carry a real suicide risk.

Ictal psychosis is a nonconvulsive status epilepticus, concurrent with the discharge and without a lucid interval, diagnosed on the electroencephalogram and aborted with an intravenous benzodiazepine; ictal depression is the seizure's own low mood, which can outlast the fit for days or longer and, rarely, prove fatal.

12 · Preictal prodrome: dysphoria and irritability before seizures

One clean discrimination and one clinical restraint carry most of the marks here. The **preictal prodrome** is the change of mood and behaviour that gathers in the run-up to a seizure and heralds it, and its examinable identity rests on a single idea: it is defined by where it sits in time, not by any symptom peculiar to it. It occupies the leading edge of the peri-ictal band, before the ictus, and belongs to the scheme that sorts the psychiatric disturbances of epilepsy by their timing to the seizure, the framework that organises the whole peri-ictal group. Two things earn easy marks. The first is that the prodrome is told apart cleanly from the aura, with which candidates perennially confuse it. The second is that a recurrent pre-seizure dysphoria is not misread as an independent psychiatric illness and treated as one.

Both errors are common and both are avoidable. The prodrome is one of the quieter members of the peri-ictal family, easy to overlook because its symptoms are ordinary, yet it rewards attention: a mood change that keeps appearing just before the fits is a datum about the epilepsy, not a separate disorder.

12.1 What the preictal prodrome is, and where it sits

A warning that comes from the seizure, before the seizure. The prodrome is the phase of altered mood and behaviour that precedes a seizure, arising from the same destabilised brain that will shortly discharge but preceding the discharge itself. Its core, as the Maudsley Prescribing Guidelines set out in their account of pre-ictal symptoms, is a **dysphoric mood** that gathers before the attack: a darkening of affect, often with tension and a shortened fuse, that the patient or those around them come to read as a sign a fit is on the way. It sits at the pre-ictal end of the axis that classifies the psychiatric disturbances of epilepsy by their timing to the seizure, the organising framework of the whole peri-ictal band, and it is that position, rather than the flavour of the mood change, that gives it its meaning. Unlike an interictal mood disorder, which lives in

the quiet baseline between attacks, the prodrome is tethered to the seizure ahead of it, and unlike the aura, it is not part of the discharge at all.

Because it is a herald rather than a disorder in its own right, the prodrome is defined less by a checklist of symptoms than by its relationship to the event it precedes. The same low, irritable state, encountered without reference to the seizures, would be worked up as a primary mood disturbance; encountered as a recurring overture to a patient's attacks, it is a prodrome, and it lifts as the seizure it announces arrives. This is why the topic is taught, above all, as a lesson in timing.

RULE **The prodrome is defined by its timing, not its symptoms.** The same dysphoria and irritability could belong to a primary mood disorder; what makes them a prodrome is that they reliably precede the patient's seizures and lift once the fit occurs. Establish the temporal relationship first, because it, and not the symptom picture, settles both the diagnosis and the first treatment move.

12.2 The clinical picture: a dysphoric, irritable build-up

Low mood and a short temper, with a scatter of somatic complaints. The prodromal state is dominated by affective and behavioural change. Kaplan and Sadock's Comprehensive Textbook of Psychiatry lists its cardinal features as **irritability** together with depression, headache and confusion, so that the picture is not purely one of low mood but of a tense, easily provoked, mildly clouded state with a somatic edge. Lishman's Organic Psychiatry describes much the same thing as **an ill-defined prodrome** of irritability, disturbed sleep, anxiety, nausea and headache appearing in the run-up to the attack, and the word ill-defined is doing real work: the prodrome rarely presents as a clean syndrome, more as a vague, uncomfortable shift that the patient struggles to name and that a spouse or parent often notices first.

- Irritability and a lowered threshold for anger, frequently the change a family reports before the patient does.
- A depressed, dysphoric mood, sometimes coloured by anxiety and inner restlessness.
- Disturbed sleep in the nights before the attack.
- Somatic complaints, with headache and nausea prominent among them.
- A degree of confusion or mental dulling, a sense of not being quite right.

The value of holding the whole cluster in mind, rather than the low mood alone, is that it keeps the clinician from mistaking the prodrome for a straightforward depressive episode. A build-up that couples irritability with disturbed sleep, headache and a vague clouding, recurring before the seizures, has a texture of its own, and recognising it is often what first raises the possibility that the mood change is seizure-related.

MNEMONIC

Long fuse, short spark

The prodrome is the **long fuse**: a non-ictal build-up of mood and behaviour that can smoulder for a variable stretch before the fit. The aura is the **short spark**: the seizure's own brief opening symptom. Holding the pair as fuse and spark fixes the two features that separate them, their length and their nature, and asserts only what this topic states.

12.3 Beyond mood: the motor prodrome

Not every warning is affective. Although dysphoria and irritability dominate the account, the prodrome is not exclusively a mood phenomenon, and one motor form is worth knowing because it is concrete and observable where the affective prodrome is vague. Lishman's Organic Psychiatry notes that in patients prone to myoclonus a **motor prodrome** may appear, the myoclonic jerks increasing in frequency over the hours leading up to a tonic-clonic seizure, so that the escalation of jerking is itself the herald of the impending convulsion. This matters clinically for two reasons. First, it is a warning the patient can often feel and a bystander can sometimes see, more reliable than the diffuse affective shift, and it therefore lends itself to precautionary action. Second, it locates the prodrome firmly within the epilepsy rather than in a psychiatric overlay: a crescendo of myoclonus is unambiguously seizure-related, and its build-up before a generalised convulsion is a reminder that the prodromal phase reflects mounting instability in a brain moving toward discharge. The affective and the motor prodrome are two faces of the same pre-ictal process, one experienced as mood, the other expressed in movement, and both are absorbed into the seizure they precede.

12.4 The temporal signature: how long before, how long lasting

The interval is wide, and its width is the point. What most sharply characterises the prodrome is not its content but its timing, and here the sources give a usable range. The Maudsley Prescribing Guidelines place the prodromal window anywhere from **thirty minutes to two or three days** before the seizure, wide enough to cover both a brief foreboding an hour before a fit and a mood that darkens across a day or two. Kaplan and Sadock's Comprehensive Textbook of Psychiatry frames the same phenomenon operationally: the prodromal symptoms begin **at least thirty minutes before** the seizure and may run on for ten minutes to three days as a continuous disturbance rather than a momentary flicker. The two accounts agree on the essentials: the prodrome opens no sooner than about half an hour before the attack and can persist for up to a few days, which is exactly the feature that separates it from the aura at one end and from an enduring interictal mood disorder at the other. A clinician who fixes this window in mind can place a reported mood change against it: a shift falling in the half-hour to a few-days range, recurring before the fits, fits the prodrome, whereas a change lasting weeks and unrelated to the attacks does not.

30 min to 2-3 days	10 min to 3 days	under 1-2 min
PRODROME PRECEDES THE SEIZURE	HOW LONG THE SYMPTOMS LAST	THE AURA, BY CONTRAST

12.5 Prodrome versus aura: the discrimination that carries the marks

Two warnings, but only one of them is the seizure. The single most examined point in this topic is the line between the prodrome and the aura, because both precede the visible convulsion and both are experienced as a warning, yet they are fundamentally different events. Lishman's Organic Psychiatry draws the distinction crisply: the prodrome is a non-ictal warning lasting minutes to days, whereas the **aura** is the initial ictal symptom, the very first fragment of the seizure itself, and is gone within **under a minute or two**. The consequences of that difference run in every direction. The aura, being the opening of the discharge, is part of the seizure and carries localising value about where the seizure begins, a matter developed with the focal seizures elsewhere in this chapter; the prodrome, being outside the discharge, localises nothing and predicts only that a seizure is coming. The accompanying figure lays out the same contrast for revision.

FEATURE	PREICTAL PRODROME	THE AURA
Position in time	a non-ictal warning before the fit	the seizure's own first symptom, at ictal onset
Duration	minutes to days	under a minute or two
Relation to the discharge	outside the electrographic seizure	the earliest part of the electrographic seizure
Typical content	a change of mood and behaviour	the seizure's opening subjective symptom
Clinical use	a warning that permits precautions	points to the seizure focus, taught with the focal seizures

The prodrome and the aura compared on the axes that separate them: the prodrome is a non-ictal warning outside the discharge, the aura the discharge's own brief first symptom, and duration together with relation to the seizure is what tells them apart.

TRAP **Calling the pre-seizure mood change an aura.** Because both precede the convulsion, candidates collapse the two, yet the aura is itself the opening of the seizure and lasts under a minute or two, while the prodrome is non-ictal and can span minutes to days. The giveaway is duration and nature: a warning measured in hours or days that lies outside the discharge is a prodrome, never an aura.

The prodrome is recognised by its timing, not its content: a non-ictal change of mood and behaviour that precedes the seizure and settles with it, unlike the aura, which is the seizure's own brief opening symptom.

12.6 Telling the prodrome from its neighbours

Several seizure-related mood states share its cross-section. The prodrome has to be separated not only from the aura but from the other affective disturbances that cluster around the seizure, and the discriminator, once again, is timing. A dysphoria that follows the fit rather than preceding it belongs to the postictal mood disturbances, the postictal depression, mania and anxiety taught in their own right elsewhere in this chapter, while a fear or affective experience that is itself the

discharge belongs to the ictal affective phenomena. A chronic, interictal irritability and dysphoria that runs in the baseline between attacks, waxing and waning without a fixed tie to any single fit, points instead to the interictal dysphoric disorder of epilepsy, a distinct entity developed separately. The prodrome is distinguished from all of these by its position: it comes before the seizure, is relieved by it, and recurs as a herald rather than persisting as a baseline state.

Hardest of all is the separation from a primary mood disorder, because the phenomenology can be identical. What settles it is the pattern over time. A depressive or irritable state that reliably appears in the day or two before seizures and remits once they occur is prodromal, however depressive it looks in cross-section; a mood disturbance that runs independently of the fits, present in the long interictal stretches and unrelieved by seizures, is a disorder in its own right and is treated as one. The two can coexist, which is why the history must map the mood onto the seizure calendar rather than assessing it in isolation.

PITFALL

A patient reports feeling low, tense and irritable for a day or two before most of their seizures, and the reflex is to diagnose a depressive disorder and start an antidepressant. When the disturbance is time-locked to the fits and lifts with them, it is the prodrome, not an independent illness, and the lever that helps is better seizure control rather than a mood drug. Treating the prodrome as a standing mood disorder both medicates the wrong target and exposes the patient to an agent whose seizure-threshold effects are weighed in their own right elsewhere in this chapter.

12.7 Why recognising the prodrome matters

Beyond the examination, the prodrome earns its keep as a clinical tool. A reliably recognised prodrome is one of the few advance warnings epilepsy offers, and it has genuine practical value. A patient who knows that a day of irritability and broken sleep tends to precede a fit can take sensible precautions, keeping away from heights, water and driving, arranging not to be alone, and ensuring any rescue medication is to hand. Families who learn the pattern can prepare rather than be caught unawares. The prodrome is also a datum for the clinician: a change in its frequency or character can signal worsening seizure control, and its consistent appearance strengthens the case that a given mood disturbance is seizure-related rather than free-standing.

The recognition cuts the other way too, protecting the patient from over-treatment. A pre-seizure dysphoria correctly identified as prodromal is not a new depressive illness and does not call for an antidepressant; naming it for what it is spares the patient an unnecessary diagnosis and an unnecessary drug. This is why the history should ask explicitly, of both patient and family, about mood in the hours to days before the attacks, and why the answer belongs in the seizure record.

GOOD TO REMEMBER

Ask every patient with epilepsy, and a family informant, specifically about a change of mood or behaviour in the hours to days before their seizures, and record it in the seizure diary. A prodrome that the patient and family learn to recognise becomes a usable warning that allows sensible precautions, and a documented pre-seizure dysphoria that clears with the fit spares the patient a mistaken diagnosis of a primary mood disorder.

IN INDIA

In much of Indian practice the family's presenting complaint is precisely the prodrome: he becomes irritable and withdrawn a day or two before the fit. Where continuous monitoring and specialist follow-up are scarce, this family-observed change is often the only warning available, and it is clinically valuable. The risk is that the pre-seizure mood change is read as a separate psychiatric illness, or attributed to non-medical causes and taken first to a faith healer, delaying the seizure-focused care that actually helps. Naming it as part of the epilepsy, and folding it into the seizure history rather than a standalone mood assessment, keeps the treatment pointed at the fits.

12.8 Management: treat the seizures the prodrome announces

The prodrome is a peri-ictal symptom, so it follows the peri-ictal rule. Because the disturbance belongs to the orbit of the seizure, its management is not a mood-disorder algorithm but the shared peri-ictal principle: it is answered first by **optimising the antiseizure medication** rather than by a psychotropic aimed at the mood. The guidelines that anchor this are, for the epilepsy itself, those of the International League Against Epilepsy (ILAE) and of NICE, with the epilepsy-specific psychiatric-management literature covering the behavioural layer; all of them frame a recurring pre-ictal disturbance as a marker that seizure control is imperfect and to be corrected there.

The **LOCUS** of care is almost always the outpatient neurology-psychiatry interface, where the regimen is reviewed and the prodrome is logged against the seizure record; it escalates to the ward or the emergency department only in the uncommon case where the prodromal state itself carries a risk of harm or heralds a severe cluster. The **FOCUS** shifts by phase: the acute target is safety through the warning window, the short-to-mid-term target is tighter seizure control, and the long-term aim is to reduce the seizures so that their heralds recede with them, while reassessing whether any mood disturbance that persists in the quiet baseline is truly interictal and deserves its own treatment. The **MODUS** spans every mode. The drug mode leads, but is pointed at the seizures: optimising the anticonvulsant is the single most useful move, and a prodrome that shortens or disappears as control improves confirms its peri-ictal nature. The psychological mode carries much of the rest, through psychoeducation, a seizure diary and teaching the patient and family to recognise the prodrome and act on it. The procedural mode enters only within the wider epilepsy workup, for the minority whose seizures resist medication. The social mode is activity and safety planning across the warning window and family involvement in spotting it. A psychotropic is reserved for a genuinely independent, interictal mood disorder, never for the prodrome itself.

13 · Postictal confusion and postictal delirium

The commonest psychiatric sequel of a seizure is also the one most often misjudged. Almost every patient who has a convulsion is briefly muddled afterwards, and for the great majority this settles without comment. **Postictal confusion** earns its place in an examination not because it is exotic but because its ordinary, self-limiting version has to be told apart from two states that share its cross-section and are anything but benign: a seizure that is still quietly running, and a psychosis that is about to erupt. It occupies the postictal slot of the axis that sorts the psychiatric disturbances of epilepsy by their timing to the seizure, the framework that organises the whole peri-ictal band, and its diagnosis is almost entirely a matter of reading that timing correctly.

The reason the marks cluster here is that the safe-looking answer is the trap. A confused, disoriented, sometimes agitated patient in the wake of a fit invites the label "postictal" and an instruction to wait, and most of the time waiting is correct. But the same picture, when it is prolonged or atypical, can be the surface expression of continuing seizure activity, and a clouded aftermath that clears and then yields to a florid disturbance is a separate syndrome altogether. The topic is taught as much by its boundaries as by its core: what postictal confusion and postictal delirium are, how the normal aftermath unfolds and clears, how seizure type and localisation reshape it, and how to distinguish it from its look-alikes. The general syndrome of delirium, its formal assessment, its subtypes and its prevention are developed in the Stroke, TBI, delirium and focal brain syndromes notes; the concern here is strictly the peri-ictal form and where it sits beside the seizure.

13.1 What postictal confusion is

A confusional state that follows the discharge and clears on its own. Postictal confusion is the clouding of consciousness that follows a seizure while the brain re-establishes its ordinary working state, the exhausted aftermath of the electrical storm rather than a disorder in its own right. Kaplan and Sadock's Comprehensive Textbook of Psychiatry frames it as a **confusional state** lasting **minutes to hours, and occasionally days**, and that single range carries most of the clinical meaning. At the short end it is the trivial, expected grogginess that needs nothing more than time and a safe place to recover; at the long end, when the state stretches towards its outer limit, it stops being reassuring and becomes a reason to look harder for something that is prolonging it. The phenomenology is that of any acute disturbance of consciousness, here generated by the seizure itself.

- Disorientation to time and place, often with a slowed, effortful grasp of what is being asked.
- Impaired attention and a reduced level of arousal, so that the patient is drowsy, distractible or difficult to rouse fully.
- Blunted, delayed responses, with fragmentary or absent recall for the period afterwards.
- Restlessness, agitation, or aimless semi-purposeful wandering in the more disturbed cases.
- A gradual, fluctuating clearing rather than an abrupt return to normal, the hallmark of a state that is resolving.

Because the disturbance is driven by a discrete, self-terminating event, its natural course is towards recovery, and that is the single most important thing to hold. It is the everyday example of the peri-ictal category, the disturbance clustered in the immediate orbit of the fit, and it is precisely its ubiquity and usually benign course that set the diagnostic trap: the clinician who has seen it resolve a hundred times assumes it always will.

13.2 When confusion becomes delirium

The same aftermath, read at greater severity. Where the postictal state is more than transient grogginess, with a clouded sensorium, prominent disorientation, perceptual disturbance and agitation, it is spoken of as **postictal delirium**. The distinction is one of degree rather than kind: both are disturbances of consciousness generated by the seizure, on a continuum from mild, brief clouding to a fuller delirious picture. What matters for the examination is not a sharp line between the two labels but that the peri-ictal form is recognised as a variant of delirium tied by its timing to a fit, and therefore expected to clear as the postictal state resolves. The formal apparatus for delirium in general, its standardised assessment, subtypes, aetiology and prevention, belongs to the Stroke, TBI, delirium and focal brain syndromes notes and is not reopened here; this topic owns only its peri-ictal expression. Naming it delirium matters because it carries delirium's cardinal feature, a globally clouded and continuously impaired consciousness, which turns out to be the key to separating it from the syndrome with which it is most dangerously confused.

13.3 The normal arc of recovery

You cannot call an aftermath prolonged without knowing the usual one. The whole clinical judgement in this area rests on an internal template of how long recovery ordinarily takes, because "prolonged" and "atypical" are defined against that baseline. Lishman's Organic Psychiatry gives the reference figure for the commonest scenario: after a tonic-clonic seizure the patient is usually alert again within **about 15 minutes**, even though a fuller sense of being back to normal, with restored energy, concentration and mood, may not return for a day or two. The gap between those two milestones is instructive. Returning to alertness is not the same as returning to baseline: a patient oriented within a quarter of an hour may still be slowed, tired and cognitively below par for a considerably longer stretch. Holding both figures keeps the clinician from two opposite errors: expecting complete restoration too soon, and, more importantly, failing to register when even the first milestone is missed. A patient who is not clearly alert well past the expected window has departed from the normal arc, and that departure is the signal to stop waiting and start investigating.

13.4 How seizure type and localisation shape the aftermath

The aftermath is patterned, not random. The duration and quality of postictal recovery track the type of seizure and the region it arose from, and these regularities are examinable because they run against intuition: one might expect a dramatic convulsion to have the shortest tail and a subtle focal event the longest, but the opposite is closer to the truth. Lishman's Organic Psychiatry records that recovery is more protracted after a generalised tonic-clonic seizure and, strikingly, surprisingly abrupt after frontal seizures, so that a patient may switch off a frontal

seizure and be back almost at once while another remains muddled far longer after a generalised fit. Localisation within the temporal lobe matters too, with prolonged postictal confusion particularly following **right temporal focal seizures**.

SEIZURE TYPE OR FOCUS	CHARACTERISTIC POSTICTAL COURSE	SOURCE
Generalised tonic-clonic	recovery more protracted than after a focal frontal event	Lishman's Organic Psychiatry
Frontal (focal)	surprisingly abrupt recovery, often almost immediate	Lishman's Organic Psychiatry
Right temporal (focal)	the substrate for a prolonged postictal confusion	Kaplan & Sadock CTP
Temporal lobe epilepsy	slow, gradual recovery with delirium unfolding over **several minutes**; postictal headache common	Lishman's Organic Psychiatry

Postictal recovery patterned by seizure type and localisation; each row is a signature the source states, and the counter-intuitive contrast between the protracted generalised aftermath and the abrupt frontal one is the examinable point.

The value of these regularities is that they let a clinician calibrate expectation to the specific seizure rather than to a single generic template. A slow, gradual clearing with a postictal headache in someone with temporal lobe epilepsy, whose anatomy and semiology are set out in their own right elsewhere in this chapter, is an unalarming pattern for that patient; the identical duration after a frontal seizure would be out of keeping and would demand explanation. The temporal and limbic structures that generate the slow, delirious, headache-laden recovery are the same ones that bind seizure activity to mood, memory and consciousness, which is why a temporal-lobe aftermath is richer and longer than the swift reset following a frontal discharge.

MNEMONIC

Frontal fast, temporal slow

The counter-intuitive recovery gradient in one line. Recovery is **surprisingly abrupt after frontal seizures**; it is **more protracted after a generalised tonic-clonic seizure**; and in **temporal lobe epilepsy it is slow and gradual over several minutes, with a postictal headache**, prolonged confusion particularly following right temporal focal seizures. Each beat asserts only a course this topic states.

13.5 Twilight states: when the boundary blurs

Neither cleanly ictal nor cleanly postictal. Not every peri-ictal disturbance resolves into a tidy discharge followed by an aftermath. A **twilight state** is the prolonged, clouded episode that arises when ictal and postictal processes run together, so that the disturbance is a protracted period of intermixed ictal and postictal changes rather than a discharge with a clean end and a recovery beginning after it. Kaplan and Sadock's Comprehensive Textbook of Psychiatry describes exactly this blurring, and it matters for two reasons. First, it explains why the neat vocabulary of the temporal axis sometimes fails at the bedside: the electrical event does not

always have a boundary that behaviour can respect, and a twilight state is the honest name for that failure of separation. Second, a state in which seizure activity and its aftermath are intermixed and prolonged is precisely when a clinician must ask whether seizure activity is still going on, which leads directly to the most important differential in this topic.

13.6 The look-alike that must not be missed: non-convulsive status

A prolonged postictal delirium and an ongoing seizure can look identical. The highest-stakes discrimination in this topic is between a genuinely prolonged postictal delirium and continuing seizure activity that never actually stopped. Lishman's Organic Psychiatry states the difficulty plainly: a prolonged postictal delirium may be very hard to distinguish from **non-convulsive status epilepticus**, because both present as a clouded, impaired consciousness in a person with epilepsy and neither announces itself with an obvious convulsion. The clinical lever the source offers is to watch for signs that the brain is still seizing: impaired consciousness accompanied by intermittent, often subtle motor signs should push the diagnosis towards ongoing status rather than a resolving aftermath. Subtle facial or eyelid twitching, rhythmic jerks, automatisms or fluctuating responsiveness in a patient who is not clearing are the features that raise the alarm. The distinction cannot always be settled clinically, and the investigation that resolves it, an electroencephalogram capturing continuing epileptiform activity, is the arbiter; the specifics of that recording belong to the EEG discussion elsewhere in this chapter. The stakes are one-directional: a resolving delirium needs only time and protection, whereas continuing status worsens the longer it runs.

TRAP **Writing off a prolonged clouded state as "just postictal".** The classic error is to attribute a persistent, fluctuating impairment of consciousness after a seizure to a slow postictal recovery and to keep waiting, thereby missing non-convulsive status epilepticus. Clinicians make it because the ordinary postictal state is so common and so reliably benign that a prolonged one is assumed to be the same thing writ large. The corrective is a rule of thumb: when a postictal state outlasts the expected arc, especially with any intermittent motor signs, treat it as possible ongoing status until an EEG says otherwise.

IN INDIA

The postictal-delirium-versus-status distinction is exactly the point at which resource constraints bite. Urgent electroencephalography, and video-EEG in particular, is concentrated in tertiary centres and is often unavailable when the decision has to be made, so the judgement frequently rests on clinical observation of consciousness and intermittent motor signs rather than on a confirmatory recording. A large epilepsy treatment gap means poorly controlled seizures, and therefore prolonged and atypical postictal states, are seen commonly. The pragmatic posture is a low threshold for suspecting continuing seizure activity in a patient who is not clearing, arranging transfer or an EEG where feasible, and not defaulting to reassurance because the confirmatory recording is out of reach.

13.7 Continuous clouding versus the lucid interval

The second boundary is drawn by the sensorium. The other syndrome that shares this territory is postictal psychosis, and here the discriminator is not severity but the continuity of

consciousness. Postictal confusion and postictal delirium cloud consciousness continuously: the patient never returns to a clear sensorium between the fit and the disturbed state, because the disturbed state simply is the unbroken aftermath of the seizure. Postictal psychosis, by contrast, declares itself only after a **lucid interval** of clear sensorium, a window of apparent full recovery before the psychosis erupts; its lucid interval, clinical features, course and management are developed in its own dedicated topic and are not repeated here. Kaufman's Clinical Neurology and Lishman's Organic Psychiatry both hang the distinction on this single hinge. The practical test is therefore about the timeline rather than the symptoms: was there a genuine return to clarity in between? If consciousness was clouded throughout, the syndrome is the confusional or delirious aftermath; if it cleared and then broke down again, it is the postictal psychosis, while the postictal affective disturbances, postictal depression, mania and anxiety, are named and taught separately. The examinable elegance is that one property, the presence or absence of an intervening clear sensorium, sorts three neighbouring postictal syndromes at a stroke.

RULE **Clouding that never clears is the signature.** Postictal confusion and delirium impair consciousness continuously from the seizure onward, with no intervening return to a clear sensorium. A state that follows a genuine lucid interval of clarity is, by definition, no longer this syndrome, and reclassifying it correctly changes the entire management plan.

13.8 Approach and management

The management is mostly good observation with a few decisive exceptions. Care of the postictal confusional and delirious state is anchored in optimising seizure control and in excluding the dangerous mimics, for which the ILAE and NICE epilepsy guidelines set the standard of care, with the epilepsy-psychiatry literature, Lishman's Organic Psychiatry among it, guiding the disturbed state itself. The governing principle is that a self-limiting aftermath should be supported and protected rather than aggressively treated, while the two look-alikes that are not self-limiting are actively sought.

The **LOCUS** of care follows severity and duration. Most postictal confusion is managed wherever the seizure was, at home, on the ward or in the emergency department, and needs only a safe, low-stimulation environment while it clears; a prolonged, atypical or non-clearing state warrants inpatient observation with access to EEG so that non-convulsive status can be excluded, and marked agitation with a risk of harm may itself justify admission. The **focus** shifts by phase: the acute target is safety together with the active exclusion of continuing seizure activity; the short-term target is supportive containment, reorientation and prevention of injury while the state resolves; and the mid-to-long-term target is improving seizure control so that the fits, and therefore their disturbed aftermaths, become less frequent. The **modus** spans every mode. The nursing and psychological mode carries most of the work: calm reassurance, a quiet setting, protection of a wandering or resistive patient, and orientation as awareness returns. The procedural mode is the EEG that settles the status question. The pharmacological mode is used sparingly: sedation is avoided in an ordinary resolving aftermath because it deepens the very clouding one is trying to read, and is reserved for confirmed continuing seizure activity, where treatment is that of the underlying status; any antipsychotic given for postictal agitation must be

weighed against its effect on the seizure threshold, a decision handled in its own dedicated topic. The social mode is family education about what a postictal state is, how long it should last and when to seek help, which also improves the collateral history on which the whole assessment depends.

PITFALL

The acutely confused postictal patient who is restless, resistive or wandering is easily met with physical restraint and heavy sedation, both of which tend to escalate injury and agitation and cloud the picture further. The confusion is transient and the patient is not wilfully uncooperative; the correct response is a calm, low-stimulation environment, gentle redirection and protection from harm while the state clears, reserving sedation for the specific situation of confirmed ongoing seizure activity.

GOOD TO REMEMBER

When a patient does not become clearly alert within the expected window after a tonic-clonic seizure, do not simply keep waiting. A postictal state that overstays its normal arc, particularly with any intermittent motor signs, should prompt an EEG to exclude non-convulsive status, and a state that follows a genuine return to clarity should redirect you towards postictal psychosis rather than a persisting delirium.

about 15 min ALERT AFTER A TONIC-CLONIC FIT	**a day or two** UNTIL FULLY RESTORED	**minutes to hours** TYPICAL CONFUSION, OCCASIONALLY DAYS	**several minutes** SLOW TEMPORAL-LOBE RECOVERY

Postictal confusion and delirium are a self-limiting, continuously clouded disturbance of consciousness in the immediate aftermath of a seizure; the two states that must never be mistaken for it are non-convulsive status epilepticus, betrayed by a non-clearing state with intermittent motor signs, and postictal psychosis, betrayed by an intervening lucid interval of clear sensorium.

14 · Postictal psychosis: lucid interval, clinical features, course and management

The mark is in the timing, not the symptom. Almost every psychotic feature of **postictal psychosis** can be produced by a dozen other conditions; what fixes the diagnosis is where the episode sits in relation to the seizure. It is a brief, **self-limiting** psychosis of abrupt onset that follows a seizure, characteristically a cluster of generalised tonic-clonic fits, and it declares itself only after a deceptive symptom-free gap. The reason it earns marks is that it is common, eminently treatable, and repeatedly misread: as a primary psychotic illness that then attracts lifelong medication, as unresolved postictal clouding, or as a psychosis arising from the seizure itself. Getting the timing right changes the management plan, the drug decision and the prognosis in a single stroke.

These peri-ictal presentations are sorted along an axis defined by their temporal relationship to the seizure, from preictal through ictal to postictal and interictal, and postictal psychosis occupies the postictal slot: it arrives after the fit but only across a delay. That single fact separates it from its two nearest mimics. It is not **ictal psychosis**, in which the psychotic phenomena are the electrographic event and there is no interval at all, and it is not the persisting **postictal confusion** of a clouded, delirious aftermath, in which the patient never returns to clarity before turning psychotic. The whole diagnosis turns on the interval, and so must the answer.

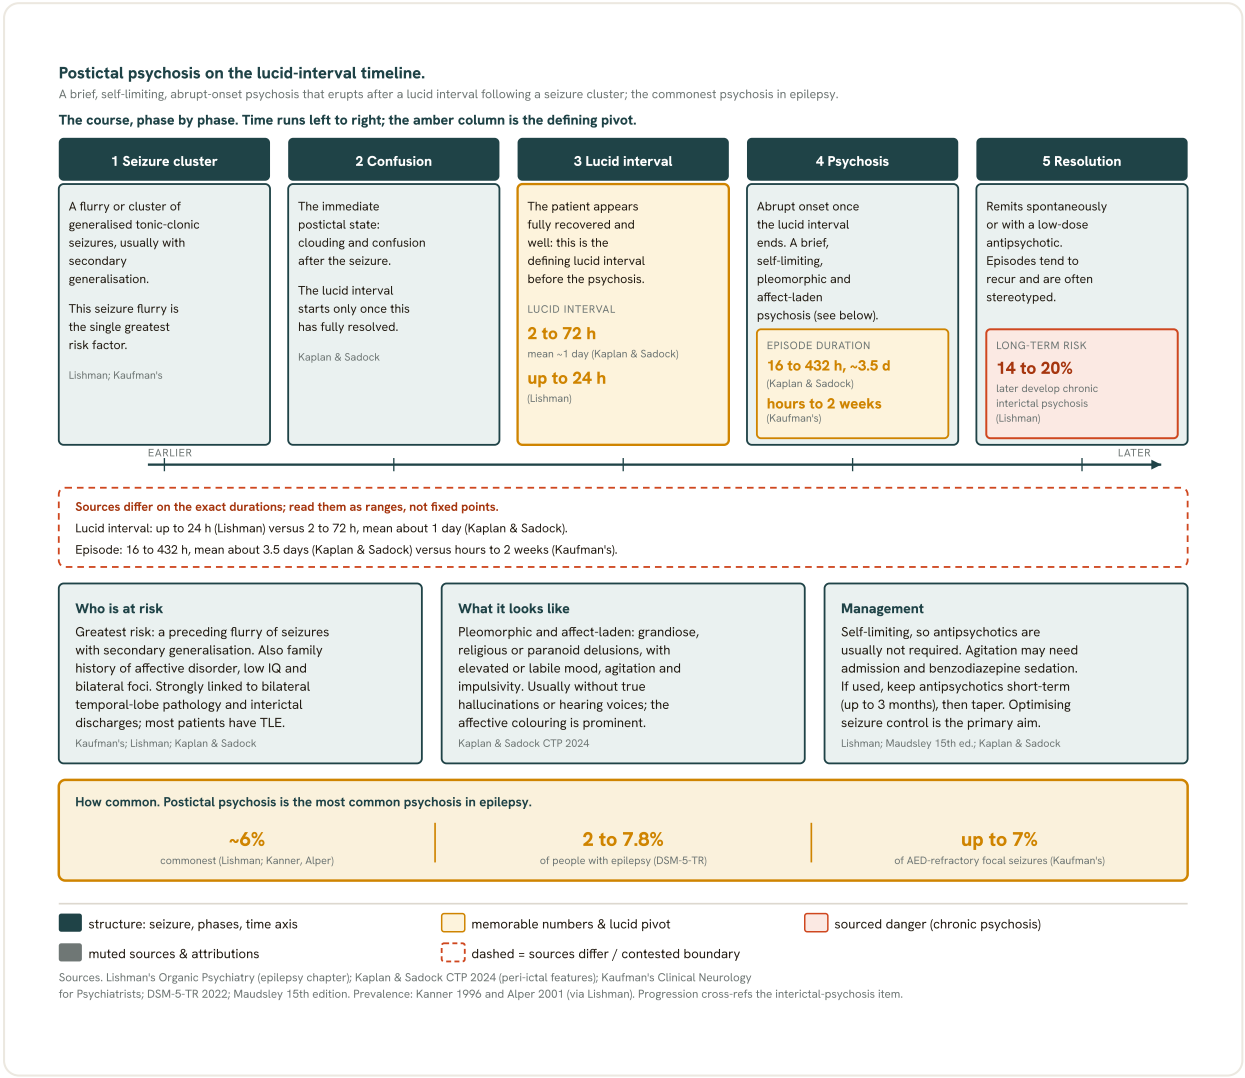


Fig 14.1 **Postictal psychosis on the lucid-interval timeline.** After a cluster of generalised tonic-clonic seizures and a brief spell of postictal confusion, the patient appears fully recovered for a defining lucid interval (2 to 72 hours, mean about 1 day in Kaplan & Sadock; up to 24 hours in Lishman) before an abrupt, brief, self-limiting and affect-laden psychosis erupts (16 to 432 hours, mean about 3.5 days; hours to 2 weeks in Kaufman's). It is the commonest psychosis in epilepsy (about 6 per cent; 2 to 7.8 per cent in DSM-5-TR), usually needs no antipsychotic, but 14 to 20 per cent later develop chronic interictal psychosis.

14.1 What postictal psychosis is

A secondary psychosis, phenotypically loud but biologically brief. Postictal psychosis is best held as a self-terminating disturbance driven by the brain state that a heavy seizure burden leaves behind, rather than as an independent psychotic disorder that merely happens to co-occur with epilepsy. Kaplan and Sadock's Comprehensive Textbook of Psychiatry and Lishman's Organic Psychiatry both describe the core picture the same way: an abrupt onset, a short course, and full

recovery, set off by seizures and usually by a run of them rather than an isolated fit. Because it is provoked by a medical event and remits when that event settles, the current classifications file it as a secondary or organic psychosis; **psychotic disorder due to another medical condition** is the DSM-5-TR home for it, and ICD-11 likewise treats the epilepsy psychoses as secondary presentations rather than as schizophrenia. Framing it this way is not pedantry. It is what justifies a short, targeted intervention and an expectation of recovery, in place of the open-ended commitment a primary psychotic diagnosis would imply. The corollary is a mechanistic one, and it is worth carrying because it explains almost every other feature of the syndrome. If the psychosis is the product of a transient, seizure-driven brain state rather than of a fixed disease process, then it should begin abruptly, run a limited course and clear as that state resolves, which is exactly what is observed. The same logic accounts for the requirement of a seizure burden, usually a cluster rather than an isolated fit, to push the brain past the threshold at which the disturbance emerges. Holding postictal psychosis as a provoked, time-limited event rather than as latent schizophrenia unmasked by epilepsy keeps the clinician from over-reading a single florid presentation.

14.2 The lucid interval: the defining feature

Recovery, then relapse. The sequence is stereotyped and is the examinable heart of the topic. The seizures end; a period of ordinary postictal confusion follows and then clears, so that the patient appears to have recovered fully; and only after that apparent recovery does the psychosis erupt. The window of apparent normality between the settling of the confusion and the onset of psychosis is the **lucid interval**. Kaplan and Sadock give it as **2 to 72 hours**, with a mean of about one day, while Lishman's Organic Psychiatry frames the same window as lasting up to 24 hours, during which the patient looks genuinely well. The interval is what makes the diagnosis feel counter-intuitive at the bedside, because the team relaxes precisely when the psychosis is about to begin. Two practical consequences follow. First, the delay means the diagnosis is frequently made in reverse, once the psychosis is under way, so the history has to be reconstructed backward to the seizure cluster and the intervening period of apparent recovery, often with a collateral informant who watched the patient look well before turning disturbed. Second, the interval is a window of opportunity as much as a diagnostic sign: a patient known to be prone to postictal psychosis, who has just had a run of fits and is currently lucid, is a patient in whom the episode can be anticipated and safety arranged before it declares itself. The gap is not empty time; it is the period in which the provoking process is still evolving toward its psychotic expression.

RULE **No lucid interval, not postictal psychosis.** The entity is defined by a symptom-free gap of hours to a few days between the seizure or cluster and the onset of psychosis. A psychosis that is continuous with the fit, or that emerges out of unbroken clouding, is a different syndrome and is managed differently.

TRAP **Collapsing the interval.** Candidates lose the mark by describing a psychosis that begins the instant the seizure stops. That erases the very feature that names the disorder and pushes the answer toward ictal psychosis or an unresolved postictal delirium. Always place the delay, and the return to lucidity, at the centre of the description.

14.3 How common it is, and the denominator problem

The commonest of the epilepsy psychoses, but the figure depends on who is counted. Postictal psychosis is the most common psychotic disorder in epilepsy, yet the reported frequency swings widely with the population sampled and the way cases are ascertained, so a single number recited without its denominator is a weak answer. The sources set out below are best read as three different questions rather than three attempts at one constant.

SOURCE	REPORTED FREQUENCY	POPULATION COUNTED
Lishman's Organic Psychiatry	about 6%	epilepsy patients in seizure-monitoring (telemetry) series
DSM-5-TR	2% to 7.8%	individuals with epilepsy
Kaufman's Clinical Neurology for Psychiatrists	**up to 7%**	drug-refractory focal seizures with impaired awareness

Reported frequency of postictal psychosis by source and denominator; the figure rises in the refractory and monitored populations, which is why the denominator must always travel with the number.

The lesson to carry into an answer is methodological. Frequencies gathered from video-telemetry units or from drug-refractory clinics are enriched for exactly the patients most prone to postictal psychosis, so they run higher than a figure drawn from an unselected epilepsy population. DSM-5-TR adds a demographic note worth holding: among older individuals the disorder is more prevalent in women. State the population, and the apparently conflicting numbers stop conflicting.

14.4 The seizure substrate and who is at risk

A bilateral, limbic disorder with a proximal trigger. Postictal psychosis is strongly associated with **bilateral temporal lobe pathology** and with bilateral, independent interictal epileptiform discharges, and most affected patients have temporal lobe epilepsy, whose anatomy and semiology are set out in their own right elsewhere in this chapter. Lishman notes that the discharges often involve the left amygdala, and, clinically important, that bilateral independent foci should be looked for and excluded before a temporal lobectomy is contemplated, because they mark a brain unlikely to be cured by removing one side; the psychiatric outcomes of epilepsy surgery are a separate topic. On this bilateral, limbic background the episode is set off by a proximal trigger, and recognising the risk profile lets you anticipate an episode before it declares itself.

- A preceding flurry or cluster of seizures with secondary generalisation, which is the **greatest single risk factor** and the usual immediate precipitant.
- A family history of affective disorder.
- Low intelligence.

- Bilateral independent seizure foci.

The cluster is doing double duty here: it is both the defining antecedent in the very definition of the syndrome and the strongest predictor that a given patient is about to develop it. A run of secondarily generalised fits in a person with bilateral temporal disease should raise the alarm before the lucid interval has even begun. The bilaterality is the thread that ties the risk profile together. A brain in which epileptiform activity is established independently on both sides, and in which the limbic structures that link seizure activity to mood and belief are implicated, is a brain more likely to generate a florid affective psychosis once a heavy seizure load is imposed on it, and less likely to be rendered seizure-free by removing a single focus. That is why the search for independent bilateral discharges is not an academic exercise but a decision that bears directly on whether surgery can help and on how closely the patient should be watched.

14.5 Clinical features: a pleomorphic, affectively coloured psychosis

Loud affect, changeable content, an affect-congruent psychosis. The picture is **pleomorphic** and carries a prominent affective colouring that is one of its most useful signatures. Kaplan and Sadock describe grandiose, religious and paranoid delusions, an elevated or labile mood, agitation and impulsivity. Hallucinations may occur, and can be prominent, but when present they tend to be affect-congruent, in keeping with the elevated or grandiose mood, rather than the sustained running commentary that marks a chronic schizophrenic illness. The mood is often high or unstable rather than flat, and the content shifts, so the cross-sectional impression can resemble mania as readily as psychosis. This affective, changeable quality, arriving abruptly after a lucid interval in a person known to have epilepsy, is what allows the syndrome to be told apart from a primary psychotic disorder; the systematic discrimination of the epilepsy psychoses from primary schizophrenia is set out in its own topic, and the phenomenology of schizophrenia itself belongs to the Schizophrenia: aetiology, phenomenology, classification and course notes. The affective loading is also why the episode is not benign in the moment. Grandiosity and religious preoccupation can carry conviction and drive action, and the mixture of elevated or unstable mood with impulsivity is the substrate for the disturbed and occasionally violent behaviour that brings these patients to attention. The content is worth taking seriously while it lasts even though the episode as a whole will pass, because the risk it generates is real during the days it runs. Reading the picture as loud but time-limited, rather than as either trivial or permanent, is the balanced clinical stance.

MNEMONIC

CLASP

The spine of postictal psychosis in five beats. **C**luster of seizures precedes it, usually secondarily generalised. **L**ucid interval before the psychosis begins. **A**ffectively coloured and pleomorphic, with grandiose, religious or paranoid delusions and any hallucinations affect-congruent. **S**elf-limiting and brief. **P**rogression, in a minority, onward to a chronic interictal psychosis. Each letter is a feature the topic actually states.

PITFALL

An abrupt, grandiose, religiously tinged psychosis with an elevated mood is easy to read as a manic episode or a first break of schizophrenia, and the reflex is to start a standard antipsychotic and plan for the long term. The affective loudness is exactly what misleads. In a patient with epilepsy who was lucid a day or two after a seizure cluster, resist the primary-psychosis label until the timing has been reconstructed.

14.6 Course and duration: self-limiting, recurrent, and the road to chronicity

Short by nature, but neither harmless nor always a single event. Left to itself the episode resolves. Kaplan and Sadock time the psychosis at **16 to 432 hours**, a mean of about 3.5 days, while Kaufman's Clinical Neurology for Psychiatrists describes episodes of hallucinations, delusions, agitation and occasional violence running from hours to two weeks. Two features of the course matter as much as the duration. First, the episodes tend to recur and are often stereotyped, so a patient who has had one is liable to have another with a similar shape, which is what makes a written relapse plan worthwhile. Second, and prognostically the sharpest point in the topic, a substantial minority do not simply recur but transform: **14% to 20%** of patients with postictal psychosis eventually develop a chronic interictal psychosis, the schizophrenia-like psychosis of epilepsy that is taught as its own entity elsewhere in this chapter. A recurring postictal psychosis is therefore not only a treatable acute problem but a marker to watch over years. The stereotypy is clinically useful in its own right: because a given patient tends to reproduce the same shape of episode, a previous presentation is a reasonable guide to the next, which lets the team recognise the pattern early and act before the disturbance peaks. The drift toward chronicity, by contrast, is the feature to fear, and it reframes what looks at first like a self-limiting nuisance into a longitudinal risk. It means that a patient with repeated postictal episodes deserves review over time rather than discharge after each one has settled, and that the goal of good seizure control is prophylactic against more than the next fit.

14.7 Management

What to actually do with the patient in front of you. Management is anchored first in optimising seizure control, for which the ILAE and NICE epilepsy guidelines set the standard of care, while the handling of the psychotic episode itself rests on the epilepsy-psychiatry literature, Lishman's Organic Psychiatry and the Maudsley Prescribing Guidelines. The organising principle is that a self-limiting condition should be treated proportionately and briefly.

The **LOCUS** of care follows the risk. The agitation, impulsivity and occasional violence of a florid episode, together with the safety concerns they carry, frequently justify inpatient admission; milder presentations can be held as outpatients with close and early follow-up, and the acutely disturbed patient may first reach an emergency setting. The **focus** shifts by phase: the acute target is agitation and behavioural risk, the short-to-mid-term reality is that the psychosis is self-limiting and needs containment rather than eradication, and the long-term aim is relapse prevention through seizure control together with vigilance for the slide into chronic psychosis. The **modus** spans every mode of treatment. Because episodes are self-limiting, an antipsychotic is

usually not required; agitation is managed with admission and benzodiazepine sedation, and where an antipsychotic is used it is given short-term, up to three months, and then tapered once the episode has resolved. Kaplan and Sadock and the Maudsley put the same point positively: the disorder remits spontaneously or responds to low-dose antipsychotics, with optimisation of seizure control as the primary aim. The choice of a particular antipsychotic against its effect on the seizure threshold is a decision handled in its own dedicated topic and is not reopened here. Alongside the drugs sit the social and procedural modes: securing the patient, educating the family, writing a relapse plan that exploits the stereotyped recurrence, and, at the level of the epilepsy itself, tightening antiepileptic treatment as the durable lever.

IN INDIA

Two realities shape this in Indian practice. The epilepsy treatment gap is large, so poorly controlled seizures and the clusters that trigger postictal psychosis are common, and the lever of seizure-control optimisation is often the single most useful thing to improve. Video-EEG telemetry is not widely available outside tertiary centres, so the diagnosis usually has to be made clinically, from a careful history of a seizure cluster followed by a lucid interval and then psychosis, rather than confirmed on a monitoring unit. The pragmatic acute plan, admission for safety with benzodiazepine sedation of agitation and only brief, cautious antipsychotic use, fits resource-limited settings well and avoids committing a patient with a self-limiting disorder to indefinite treatment.

GOOD TO REMEMBER

When a person with epilepsy becomes psychotic a day or two after a cluster of seizures, having clearly been lucid in between, think postictal psychosis first. Admit if there is agitation or risk, sedate with a benzodiazepine, optimise seizure control, use any antipsychotic at low dose and for a short time, and do not reflexively begin a lifelong antipsychotic for what is, by nature, a self-limiting episode.

14.8 Telling it apart from its neighbours

The same axis that classifies it also distinguishes it. Because postictal psychosis is defined by its position after a lucid interval, the cleanest way to separate it from the conditions it is confused with is to walk the timeline. It is not the psychosis that coincides with the electrographic seizure and has no interval, which is the ictal form. It is not the clouded, disoriented aftermath that never clears into lucidity, which is postictal confusion and delirium, held as its own peri-ictal syndrome. And it is not the persistent, hallucination-rich psychosis that runs on between seizures independent of them, which is the chronic interictal or schizophrenia-like psychosis, the very endpoint that a minority of postictal cases eventually reach. The formal separation of any of these from a primary schizophrenic illness, resting on preserved affect, intact premorbid function and the tight relation to seizures, is developed in the dedicated discrimination topic, and the features of primary schizophrenia sit in the Schizophrenia: aetiology, phenomenology, classification and course notes. Keep the axis in mind and the four pictures stop blurring into one another.

2 to 72 h	16 to 432 h	about 6%	14 to 20%
LUCID INTERVAL BEFORE ONSET	EPISODE DURATION	COMMONEST EPILEPSY PSYCHOSIS	PROGRESS TO CHRONIC PSYCHOSIS

Postictal psychosis is a brief, self-limiting, affectively coloured psychosis that erupts after a lucid interval of hours to a few days following a cluster of seizures; it is usually managed with containment and short-term treatment rather than long-term antipsychotics, but a minority go on to a chronic interictal psychosis.

15 · Postictal depression, postictal mania and postictal anxiety

The trap is to treat the mood when the task is to treat the seizure. When a person with epilepsy turns low, anxious or, far less often, elated in the days after a fit, the cross-sectional picture can be indistinguishable from a primary mood or anxiety disorder, and the reflex to reach for an antidepressant is strong. What separates these postictal affective disturbances from the free-standing illnesses they imitate is not their phenomenology but their position in time: they arrive tethered to the seizure, and the depressive ones in particular behave quite differently under treatment. The marks are earned by holding two facts together - that mood and anxiety symptoms clustered after a seizure are common and real, and that the treatment logic for them is inverted relative to what the symptoms alone would suggest.

These presentations are located on the axis that sorts every peri-ictal disturbance by its relationship to the seizure, from preictal through ictal to postictal and interictal. That classifying framework is developed as its own topic and is borrowed here only to place the syndromes at hand; all occupy the postictal slot, the disturbed brain state that follows the discharge. This account works through the three named members in turn - the common depressive picture, the sparser anxiety one and the rare manic one - then turns to the lateralisation question, the postictal flare of a settled interictal disorder, and the management principle that gives the topic its examination edge.

15.1 The postictal window: when these symptoms appear

Fix the timing first, because the timing is the diagnosis. The defining fact about this group is that the symptoms sit in a band after the seizure rather than floating free of it. The Maudsley Prescribing Guidelines place the postictal psychiatric symptoms - depression, anxiety, suicidal ideation and psychosis together - in a window of **several hours to 7 days** after the fit. That window does two things. It tells you where to look, so that a new low mood or a surge of anxiety in the days after a seizure should prompt the timing question before a prescription; and it tells you what company these symptoms keep, since they travel as a recognisable postictal cluster rather than as isolated events. Two members of that cluster are taught in full elsewhere and are only named here: the suicidal ideation that belongs to the account of suicide risk in epilepsy, and the postictal psychosis that declares itself after a lucid interval. The window also warns against symmetry: a disturbance that begins during the discharge is ictal, and one that runs on independently in the baseline is interictal, so a mood change dated to the days after a fit is what earns the postictal label.

15.2 Postictal depression: the common one

The everyday member of the group, and the one most often misattributed. Of the three affective pictures, **postictal depression** is the one the psychiatrist will actually meet. Kaplan and Sadock's Comprehensive Textbook of Psychiatry describes it as common and makes the clinically important point that a prolonged depressive state may follow focal impaired-awareness seizures even when the seizure itself carried no depressive content. That is the subtle observation worth carrying: it is tempting to expect a depressive aftermath only from a seizure with an unpleasant affective aura, but the low state can settle in after a fit that was, in its own experience, affectively neutral. The depression is a property of the postictal brain state rather than a carry-over of the ictal mood, and can outlast the ordinary confusion of the immediate aftermath. The label must be handled with discipline, because two other depressions in epilepsy are easily confused with it and owned by their own topics: the ictal depressive phenomena that are the psychiatric face of the discharge, and the interictal depression of epilepsy, whose prevalence, phenomenology and screening are set out separately and which behaves very differently under treatment. Naming the postictal form precisely changes the expectation of course and response: it is a self-limiting state locked to a recent seizure, not a free-standing depressive episode.

15.3 Postictal anxiety: real but thinly described

Named in the cluster, but the evidence is sparse. **Postictal anxiety** is reported as part of the postictal symptom cluster, alongside the depressive, suicidal and psychotic elements that share the same window. Beyond that listing the standard corpus says little: its phenomenology, frequency and natural history are not worked out to anything like the depth of the depressive picture, and an honest answer acknowledges that thinness rather than papering over it. What can be said with confidence is where postictal anxiety sits relative to its neighbours. It is not the brief, stereotyped ictal fear or ictal panic that is the discharge itself, taken up among the ictal affective phenomena, and it is not one of the free-standing anxiety disorders of epilepsy that run in the interictal baseline and are treated in their own right. It is a postictal phenomenon: an anxious colouring to the aftermath of a seizure, dated to the same hours-to-days band as the rest of the cluster. The consequence is modest. Recognise postictal anxiety as a legitimate member of the group, keep it separate from ictal panic and from an interictal anxiety disorder on grounds of timing, and resist importing a management or prognosis the literature has not established.

15.4 Postictal and epilepsy-related mania: the rare one

The exception, and examinable precisely because it is rare. If depression is the rule, **postictal mania** is the exception. Kaplan and Sadock's Comprehensive Textbook of Psychiatry records that manic or mixed features are much rarer than depressive features in the **mood disorder due to epilepsy** picture, so an elevated or expansive postictal mood is a genuine but uncommon event. Its rarity is the very thing that makes it a favourite discriminator, because a manic presentation in a person with epilepsy pulls the differential in several directions at once. The timing of its emergence is itself informative and slightly paradoxical: the same source notes that manic symptoms may rarely emerge either with an increase in seizure frequency or, conversely, after seizures come under control. The second route is the counter-intuitive one and deserves care,

because a mood disturbance surfacing as the fits are suppressed echoes the pattern named as forced normalisation (the Landolt phenomenon) and alternative psychosis, which is developed as its own topic and only pointed to here. The lesson is that neither a worsening nor an improvement in seizure control excludes an elevated mood; the manic picture, though rare, sits at both ends of the seizure-control spectrum rather than one predictable point.

MNEMONIC

Common down, rare up, date it to the fit

The postictal affective triad in one line. **Down** is common: postictal depression is the member you will actually see. **Up** is rare: manic or mixed features are far less frequent than depressive ones. **Anxiety** travels in the same cluster. And every one of them is **dated to the fit**, appearing in the window of hours to days that gives the group its name. Each beat asserts only what this topic states.

15.5 The lateralisation question in postictal mania

A tentative anatomy, offered with its uncertainty intact. When mania does arise in epilepsy, the natural question is whether it maps onto a particular seizure focus, and the honest answer is qualified. Kaplan and Sadock's Comprehensive Textbook of Psychiatry names a **right temporal focus** as a possible source of mania in epilepsy while being explicit that the association is not well established. That qualification is not a footnote to be dropped; it is the substance of the point. The lateralisation fits a broader intuition that right-sided limbic disturbance tends toward elevated or disinhibited states and left-sided toward the depressive, but the epilepsy data supporting it are thin. The temporal lobe epilepsy that provides the substrate has its anatomy and semiology set out in their own topic and is not re-taught. What this fragment owns is the narrow, hedged proposition itself: a right temporal focus is a candidate source of epilepsy-related mania, no more firmly than that. Carrying the uncertainty forward is what separates a good answer from an overconfident one.

TRAP **Stating the right-temporal link as established fact.** Candidates read that mania in epilepsy has been associated with a right temporal focus and then present it as a firm localising rule. The source flags the association as not well established, and the mark is in reproducing that hedge. Assert the lateralisation as possible and unconfirmed, not as a law.

15.6 Postictal breakthrough of a settled interictal disorder

A separate mechanism hiding inside the same time window. Not every postictal disturbance is a fresh, self-contained syndrome. The Maudsley Prescribing Guidelines note that patients who carry an interictal psychiatric disorder may experience a postictal worsening of previously-remitted symptoms, best held as **breakthrough symptoms** rather than as a new illness. The distinction is mechanistically real and clinically consequential. In one patient the postictal state generates a de novo depressive or anxious picture where there was no standing disorder; in another, the same state briefly reactivates a mood or anxiety disorder that had been in remission, so that what looks like a relapse is really a transient, seizure-triggered flare on the back of a recent fit. Reading the second as the first is where clinicians go wrong, and the error has a direction: it

tends toward over-treatment, because a brief reactivation of a settled disorder is easily mistaken for a durable relapse that then attracts a permanent escalation of maintenance medication. The corrective is to date the worsening to the seizure and watch it settle as the postictal state resolves, rather than to rewrite the long-term plan on the strength of a few post-seizure days. Breakthrough symptoms are the postictal window acting on an already-vulnerable brain, and belong in the differential of any apparent relapse that follows a fit.

PITFALL

A patient whose depression or anxiety had settled becomes symptomatic for a few days after a seizure, and the team reads it as a relapse of the underlying disorder and steps up the maintenance antidepressant or adds an agent for the long term. If the worsening is dated to the fit and clears as the postictal state resolves, it is a breakthrough phenomenon, not a true relapse, and a permanent change of regimen is the wrong response to a transient, seizure-driven flare.

POSTICTAL AFFECTIVE SYNDROME	FREQUENCY	CHARACTERISTIC RELATION TO THE SEIZURE	THE POINT THAT EARNS THE MARK
Postictal depression	common	may be prolonged, and can follow focal impaired-awareness seizures even when the ictus was not depressive	does not respond to antidepressants; the lever is seizure control
Postictal anxiety	reported in the cluster	shares the postictal window with depression, suicidal ideation and psychosis	corpus detail is sparse; separate it from ictal panic and from an interictal anxiety disorder
Postictal / epilepsy-related mania	rare	emerges rarely with rising seizure frequency, or after seizure control	a right temporal focus is a possible but unconfirmed source

The three postictal affective syndromes compared on frequency, relation to the seizure and examinable point; the suicidal and psychotic members of the same cluster are developed in their own topics and named here only to locate them.

15.7 Management: the antidepressant paradox

Here the topic pays off, because the obvious treatment is the wrong one. The governing principle of this group is a negative one, and it is the single most examinable fact in the fragment. The Maudsley Prescribing Guidelines state that peri-ictal depressive symptoms do not appear to respond to antidepressants and are managed instead by optimising the antiseizure medication. This inverts the instinct that a depressive picture calls for an antidepressant, and it draws the line between this topic and the interictal depression of epilepsy, set out in its own account, which does respond to antidepressant treatment. The same contrast holds against the interictal dysphoric disorder of epilepsy, the chronic baseline mood picture taught as its own entity, which likewise answers to an antidepressant. A **peri-ictal depression** tied to recent seizures is the

exception, and treating it as an ordinary depression wastes weeks on a drug that will not work while the real lever goes untouched.

The structure of care follows from that principle. The **LOCUS** of management is overwhelmingly the outpatient epilepsy or neuropsychiatry clinic, where seizure control is reviewed and adjusted; inpatient care is reserved for the few in whom postictal risk, particularly the suicidal ideation that shares the cluster, is acute, and the emergency department enters only when safety cannot be held otherwise. The **focus** shifts by phase: the acute target is safety and containment across the few postictal days, the short-to-mid-term aim is to let a self-limiting state resolve without committing the patient to a mood-specific drug the evidence says will not help, and the long-term aim is to reduce the seizure burden that keeps generating the disturbances. The **modus** is weighted away from psychotropic prescribing and toward the epilepsy itself: the primary mode is optimisation of the antiseizure regimen, supported by educating patient and family that these symptoms are seizure-driven and time-limited, with psychological support and monitoring across the postictal window rather than a reflexive antidepressant. The guideline anchor is the epilepsy standard of care set by the ILAE and NICE, within which seizure optimisation is defined, while the handling of the psychiatric symptoms rests on the epilepsy-specific psychiatric-management literature rather than the ordinary mood-disorder algorithms.

RULE **Peri-ictal depressive symptoms are treated by optimising seizure control, not by antidepressants.** They do not appear to respond to antidepressant drugs and are managed instead by optimising the antiseizure medication. This is the exact reverse of the interictal depression of epilepsy, and the reversal is the whole clinical point.

IN INDIA

The treatment gap in Indian epilepsy care means that poorly controlled seizures, and the postictal mood and anxiety disturbances they generate, are common in ordinary practice. That reality aligns with the management principle rather than fighting it: because the response to peri-ictal depression is to optimise seizure control rather than add an antidepressant, the most useful and affordable intervention is often better antiseizure treatment with the medicines already in use, not a new psychotropic. Where telemetry is unavailable the timing must be reconstructed from a careful history - a fit, then low or anxious mood over the following days - so a seizure-driven picture is not mislabelled as a primary depression and started on a drug that will not touch it.

GOOD TO REMEMBER

When a person with epilepsy becomes depressed or anxious in the days after a seizure, date the symptoms to the fit before you prescribe. If they sit in the postictal window, the first move is to review and optimise seizure control, not to start an antidepressant, and to watch a self-limiting state settle. Keep the suicidal-ideation risk of the postictal cluster in mind while you wait.

15.8 Holding the group apart from its neighbours

The same timing axis that groups these syndromes also separates them from their mimics. Because the postictal affective disturbances are defined by their band after the seizure, the

cleanest separation is to walk the timeline. They are not the ictal affective phenomena, the fear, panic or low mood that is the discharge itself and clears when the seizure ends. They are not the interictal mood and anxiety disorders of epilepsy, which run in the baseline between attacks and, crucially, do respond to the usual treatments; the interictal depression of epilepsy and the interictal dysphoric disorder both sit on that side of the line. And within the postictal band, the mood and anxiety members must be kept apart from postictal confusion and delirium, a disturbance of arousal and orientation rather than mood, and from the postictal psychosis that emerges after a lucid interval, each a peri-ictal syndrome in its own right. The discipline is the same: establish where the disturbance sits relative to the seizure, and let that placement, not the surface symptom, decide both the label and the first treatment move.

- Depressive, anxious and, rarely, manic in colour, all appearing in the postictal window rather than during the discharge or in the baseline.
- Set apart from the ictal affective phenomena by the interval after the seizure, and from the interictal mood and anxiety disorders by their tight tie to a recent fit.
- Set apart from postictal confusion and delirium, which disturb arousal and orientation, and from postictal psychosis, which follows a lucid interval.
- Managed, in the depressive case, by optimising seizure control rather than antidepressants, the reverse of the interictal picture.

Common	Rare	Up to 7 days	No AD response
POSTICTAL DEPRESSION	POSTICTAL / EPILEPSY MANIA	POSTICTAL SYMPTOM WINDOW	PERI-ICTAL DEPRESSION

Postictal depression, anxiety and the rare postictal mania appear in a window of hours to a week after a seizure; the depressive picture is common and, unlike interictal depression, does not respond to antidepressants but to optimisation of seizure control, while manic or mixed features are rare and only tentatively linked to a right temporal focus.

16 · Interictal dysphoric disorder of epilepsy

The commonest way mood disorder shows itself in epilepsy is not the one candidates are trained to look for. Faced with a person who has epilepsy and is clearly unwell in mood, the reflex is to reach for the template of a major depressive episode and test the history against it. That template usually fails, because the disturbance most such patients carry is a chronic, shifting, irritable low mood that never resolves into a clean depressive episode and never fully lifts either. Blumer gave this picture a name and a place of its own, and thought it central enough to call it **the major psychiatric disorder of epilepsy**. It earns marks precisely because it is easy to miss: it hides behind a search for classic depression, borrows features from the anxiety and even the psychotic disorders, and responds well to treatment once it is recognised for what it is.

Everything about the syndrome follows from where it sits in time. The psychiatric disturbances of epilepsy are sorted along an axis running from before the seizure, through the discharge, to its

aftermath and finally to the quiet baseline between attacks; that classification is developed in its own right elsewhere in this chapter and only borrowed here. Interictal dysphoric disorder belongs to the last position. It is an interictal disorder that runs independently of the timing of seizures, which separates it from the peri-ictal mood changes, the preictal prodrome of dysphoria and irritability and the postictal depression, mania and anxiety, that keep time with the fit and are taught as their own topics. The mood change here neither dates to the last seizure nor predicts the next; it occupies the background of the illness, not its edges.

16.1 The disorder and its two names

One entity under two labels, and a deliberate claim about its weight. The syndrome entered psychiatry through Blumer's account, published in **1991**, of a mood disturbance specific to epilepsy that did not fit the depressive categories then in use; he named it the **interictal dysphoric disorder**. More than a decade later, in **2003**, **Kanner** described the same picture as the **dysthymia-like disorder of epilepsy**, a label that foregrounds its resemblance to a chronic low-grade depression rather than to an acute episode. Both names point at one referent and both are worth carrying, because an examiner may use either. Blumer went further than merely naming it. Reviewing the psychiatric burden of epilepsy in Kaplan and Sadock's Comprehensive Textbook of Psychiatry, he characterised this condition as the major psychiatric disorder of the illness, a phrase that stakes a claim: this chronic, pleomorphic dysphoria is the one the epilepsy clinician meets most often and should know best. It is a clinically described syndrome rather than a category with its own operational criteria in the standard manuals, which is part of why it slips past a diagnostic process built around them. Naming it correctly is therefore the first clinical act, because a disorder with no line in the checklist is a disorder that goes untreated.

16.2 Interictal by definition: why the timing decides everything

The name locates the disorder before it describes it. The word interictal is doing diagnostic work. A peri-ictal mood change announces its own origin by its timing, so a dysphoria clustering in the hours around a fit is close to self-explanatory. This disorder does the opposite. Both Lishman's Organic Psychiatry and Kaplan and Sadock's Comprehensive Textbook of Psychiatry place it in the interictal baseline, arising independently of seizure timing rather than as a peri-ictal phenomenon, so it neither dates to the last seizure nor forecasts the next. That has two consequences the axis is designed to make obvious.

- It must be assessed as a mood disorder in its own right, with a full mood and risk history, rather than explained away as a transient reaction to the last seizure.
- Its treatment is aimed at the mood disorder itself and does not follow automatically from better seizure control, which is the primary lever for the peri-ictal syndromes but not for this one.

Placed rightly, the disorder is held as both truths at once: a genuine, treatable mood disorder that nonetheless belongs to the epilepsy, so that it is neither dismissed as a passing aftermath of the fits nor cut loose from its specific tie to the illness.

RULE **The disorder keeps no appointment with the seizure.** Interictal dysphoric disorder runs in the baseline state and does not track the fit, so it is diagnosed and treated as a mood disorder in its own right, with the epilepsy as vulnerability and context rather than as the proximate trigger. A dysphoria that clearly dates to the hours around a seizure is a peri-ictal syndrome and belongs to a different topic.

16.3 The core picture: a chronic dysthymia broken by brief normal mood

Chronic, but never steady. Lishman's Organic Psychiatry gives the central shape of the syndrome plainly: a **chronic dysthymia** interrupted at frequent intervals by brief periods of normal mood. Two words carry the diagnosis. It is chronic, so the low mood is the background state of the illness, present far more often than not and measured over months and years, not as an episode with a beginning and an end. And it is interrupted, so the depression is not sustained: short intervals of feeling well break it up, and those windows make the course intermittent and labile rather than a single depressive episode. The clinical lesson is that the disorder is defined as much by its rhythm as by its content. A patient seen in a good-mood window can look euthymic, and an assessment that lands in a normal-mood window will underestimate or miss the disorder. The waxing and waning is intrinsic and must not be read as a partial treatment response or a stable recovery. This fluctuating, low-grade, chronic quality is precisely why it does not answer to the description of a classic major depressive episode, whose prevalence, phenomenology and screening in epilepsy are set out as their own topic elsewhere in this chapter.

16.4 Paroxysmal irritability and agitation

Not sadness alone: the irritable, agitated edge. Kaplan and Sadock's Comprehensive Textbook of Psychiatry emphasises that the disorder is closely associated with **paroxysmal irritability** and agitation, and this is one of the features that pulls it furthest from ordinary depression. The irritability is not a steady, low-level bad temper running through the illness; it comes in paroxysms, in bursts that are sudden, disproportionate to their trigger and short-lived, with agitation alongside. The irritable component is often what the family reports, rather than the low mood the patient may never volunteer, so the complaint can be friction at home rather than sadness. It also feeds the diagnostic confusion the disorder is notorious for: the irritability can be read as a personality problem or as behavioural disturbance, each reading missing the mood disorder underneath. The irritability belongs to the dysphoric disorder as one facet of a single condition, not a separate diagnosis bolted on. Sustained interictal irritability, aggression and episodic dyscontrol in epilepsy form a behavioural syndrome with its own dedicated topic; here the irritability is one colour of the mood disturbance, arriving in paroxysms against the chronic dysthymic background rather than as a fixed trait.

16.5 The reach toward psychosis: paranoia and hallucinations

The upper border of the syndrome touches psychosis. At its more severe pole the disorder does not stay purely affective. Kaplan and Sadock's Comprehensive Textbook of Psychiatry notes that it may carry accompanying **paranoia and hallucinations**, which places it on a continuum with the psychotic disorders rather than in a sealed mood compartment. Mood and psychosis in epilepsy are not cleanly separable, and the same patient's dysphoria can shade into paranoid

ideas and perceptual disturbance without becoming a different illness. The continuum explains why a dysphoric epilepsy patient can present with a paranoid, quasi-psychotic colouring that would baffle anyone holding the disorder to be a simple depression. It should not, however, be over-read. Some paranoia or hallucination at the severe end of the dysphoric spectrum does not convert the picture into a primary psychotic illness, and the chronic interictal psychosis of epilepsy, the schizophrenia-like psychosis, is a separate entity taught in its own right. The task is to recognise the psychotic tinge as the far edge of the mood disorder while keeping the frank, persistent psychoses of epilepsy conceptually distinct.

PITFALL

When the dysphoria arrives with paranoid ideas or a hallucination, the reflex is to read it as a primary psychosis and start an antipsychotic as the main treatment. That inverts the therapeutic logic. The condition is fundamentally a mood disorder that responds to antidepressant treatment, with a low-dose neuroleptic reserved as an addition where the psychotic colouring is prominent, not as the first or the sole agent. Reaching for antipsychotic monotherapy leaves the mood core untreated.

MNEMONIC

DIP

The three domains of the disorder held in one word. **D**ysthymia that is chronic and broken only by brief windows of normal mood. **I**rritability, paroxysmal, with agitation. **P**aranoia and hallucinations at the severe pole, marking the continuum with psychosis. Each letter is a feature the disorder actually shows.

16.6 Reading the domains together: a pleomorphic disorder

Several kinds of symptom, one condition. The reason interictal dysphoric disorder resists the single-episode template is that it spreads across more than one symptom domain at once. It is depressive, in its chronic, fluctuating low mood; irritable and agitated, in the paroxysms that punctuate that background; and psychotic at its edge, where paranoia and hallucinations can appear. This pleomorphism is not vagueness but the defining texture of the syndrome, and it is exactly what a checklist built for a discrete depressive episode cannot capture. The accompanying figure lays the domains out together; the table below pairs each with its character and the misreading it invites.

SYMPTOM DOMAIN	CHARACTER IN THE DISORDER	HOW IT IS MISREAD
Depressive	chronic dysthymia broken by brief windows of normal mood	a good-mood window is mistaken for recovery or a stable baseline
Irritable	paroxysmal irritability with agitation	read as a personality trait or isolated behavioural disturbance
Psychotic-adjacent	possible paranoia and hallucinations	mistaken for a primary psychotic illness

The domains of interictal dysphoric disorder, each with the form it takes in the syndrome and the diagnostic error it invites; the disorder is a single, pleomorphic condition, not three separate diagnoses.

Filed separately, the low mood becomes a depression, the irritability a personality issue and the paranoia an unrelated psychosis, and no one treats the syndrome that unites them.

16.7 The commonest face of mood disorder in epilepsy

This, not the classic episode, is what you will meet most. The single fact that gives the topic its clinical weight is one of base rates. Kaplan and Sadock's Comprehensive Textbook of Psychiatry states that of patients who have both epilepsy and a mood disorder, most carry this dysthymia-like, interictal picture rather than a classic major depressive disorder. In the general population the prototype of a mood disorder is the depressive episode, but in the epilepsy clinic the prototype is this chronic, fluctuating, irritable dysphoria, and the clean depressive episode is the less typical presentation. The consequence is unforgiving: screening only with major-depressive-episode criteria systematically under-detects the very disorder most likely to be present. The prevalence, phenomenology and screening of depression in epilepsy are developed as their own topic; the point here is the epidemiological weighting, that when mood disorder appears in epilepsy it wears this face more often than the textbook one, so the assessment must catch a pattern rather than confirm an episode.

TRAP **Screening only for the classic depressive episode.** Candidates lose the mark, and clinicians miss the diagnosis, by testing the epilepsy patient against major depressive disorder, finding no clean episode, and concluding the mood is normal. The commonest mood presentation in epilepsy is this dysthymia-like disorder, which will not satisfy those criteria. Look for a chronic, intermittent, irritable dysphoria, not a discrete episode.

16.8 The seizure substrate and localisation

An association with frequent focal seizures, tentatively lateralised. The disorder is not evenly distributed across everyone with epilepsy. Kaplan and Sadock's Comprehensive Textbook of Psychiatry links interictal dysphoria with **frequent focal impaired-awareness seizures**, and, more tentatively, with a possible preponderance of left temporal foci, while being explicit that the lateralisation is not well established. Both halves of that claim matter, and the second must keep its caution. The association fits the wider impression that a heavier, poorly controlled seizure burden of this type carries a greater psychiatric load; the distinction between focal aware and focal impaired-awareness seizures, and the anatomy and semiology of temporal lobe epilepsy behind it, are set out in their own topics and only named here. The suggestion of a left temporal preponderance is offered as a possibility rather than a fact, and should be stated that way in an answer; asserting a firm lateralisation overstates what the source supports. The neurobiology linking the limbic structures of the temporal lobe to mood and behaviour is the chapter's mechanism topic, and it supplies the plausible substrate: an interictal disorder commoner where focal impaired-awareness seizures are frequent fits a mood disturbance generated by chronically disordered limbic function, even where the side of the focus stays uncertain.

16.9 Management

The recognition is most of the battle, but the treatment is genuinely rewarding. The management question is anchored, for the epilepsy itself, in the International League Against Epilepsy and the National Institute for Health and Care Excellence epilepsy guidelines, and, for the mood disorder, in the epilepsy-specific psychiatric-management literature described by Kaplan and Sadock's Comprehensive Textbook of Psychiatry and Lishman's Organic Psychiatry. Encouragingly, this is a treatable disorder that, once named, usually responds well.

The **LOCUS** of care is, for the great majority, the outpatient clinic, where a chronic and fluctuating mood disorder is best followed over time; it escalates to inpatient care only where severity, prominent paranoia or hallucinations, or risk make it necessary, and an acutely agitated or psychotic presentation may reach an emergency setting. The **focus** shifts by phase: the acute target is the dysphoria, the paroxysmal irritability and any distressing psychotic colouring; the mid-term target, given the chronic relapsing course, is sustained stabilisation of mood rather than the resolution of a single episode; and the long-term aim is maintenance and vigilance, because a disorder that waxes and wanes over years is one to hold in view rather than discharge after a good spell. The **modus** spans the available modes of treatment.

- The pharmacological mode is led by an **antidepressant**, to which the disorder shows a good response, with a low-dose neuroleptic added where the picture reaches toward paranoia or hallucinations; the choice of a particular antidepressant against its effect on the seizure threshold is handled in its own dedicated topic and not reopened here.
- The psychological and social modes carry real weight: explaining to patient and family that the chronic, fluctuating, irritable low mood is part of the epilepsy and is treatable reframes what has often been read as difficult behaviour, and support is needed across a long course.
- Attention to the epilepsy itself stays in the background: the disorder is commoner where focal impaired-awareness seizures are frequent, so good seizure control remains part of good care, though, this being interictal, it is not the primary lever it is for the peri-ictal syndromes.

IN INDIA

Two realities shape this in Indian practice. First, the epilepsy consultation is dominated by seizure control, often against a wide treatment gap and little time per patient, so a chronic, fluctuating, irritable dysphoria the patient never volunteers is easily overlooked; actively asking about mood and irritability, and asking the family, is what surfaces it. Second, a widespread and understandable hesitancy about prescribing antidepressants in epilepsy, for fear of provoking seizures, can leave an eminently treatable mood disorder untreated. Since the disorder responds well to antidepressants, the task is to treat it while attending to seizure-threshold safety through the dedicated topic on antidepressant choice, not to withhold treatment by default.

GOOD TO REMEMBER

When you assess a person with epilepsy, screen for a chronic, intermittent, irritable low mood, not only a classic depressive episode, and ask the family about paroxysmal irritability the patient may not report. If you find it, treat it: the disorder responds well to an antidepressant, with a low-dose neuroleptic added only where paranoia or hallucinations are prominent. Naming the pattern is the step that unlocks a good outcome.

Interictal	**Dysthymia-like**	**Commonest**	**Antidepressant**
RUNS FREE OF SEIZURE TIMING	CHRONIC, FLUCTUATING, IRRITABLE MOOD	MOOD PRESENTATION IN EPILEPSY	GOOD TREATMENT RESPONSE

Interictal dysphoric disorder is a chronic, fluctuating, irritable dysphoria that runs in the baseline between seizures, spreads across depressive, irritable and mildly psychotic domains, is the commonest mood presentation in epilepsy rather than a classic depressive episode, and responds well to an antidepressant, with a low-dose neuroleptic added only where paranoia or hallucinations appear.

The psychoses of epilepsy

17 · Schizophrenia-like psychosis of epilepsy (chronic interictal psychosis): features and management

This is the epilepsy psychosis that most faithfully counterfeits schizophrenia, and recognising the counterfeit for what it is carries the marks. The **chronic interictal psychosis** of epilepsy, loosely called the **schizophrenia-like psychosis of epilepsy**, is a persistent psychotic disorder that arises between seizures, in **clear consciousness**, in a person with long-standing epilepsy, and that reproduces the surface of schizophrenia closely enough to be mistaken for it on a single mental state examination. The reason it is examined so reliably is that it forces the candidate to hold two truths at once: it looks like schizophrenia, and yet it is a secondary psychosis with its own age of onset, its own temporal relationship to the epilepsy, its own risk profile and its own course. A trainee who describes only the resemblance has answered half the question; the marks are in the resemblance *and* the signature that betrays it.

The psychoses of epilepsy are sorted along an axis defined by their timing to the seizure, running from preictal through ictal to postictal and finally to the interictal, seizure-independent forms. This entity occupies the interictal, chronic slot: it runs on independently of individual seizures and of seizure control rather than being locked to a fit. That single fact separates it from its seizure-bound neighbours, each of which is taught in its own right in this chapter. It is not the ictal psychosis that is itself the electrographic discharge, it is not the brief postictal psychosis that erupts after a lucid interval following a cluster of fits, and it is not the alternative psychosis of forced normalisation, in which the disturbance appears as the electroencephalogram paradoxically normalises. The systematic discrimination of this psychosis from primary

schizophrenia is developed in its own dedicated topic, and the phenomenology of the primary illness itself belongs to the Schizophrenia: aetiology, phenomenology, classification and course notes; here the task is the epileptic entity itself, its features and its management.

17.1 What the entity is

A chronic, interictal, schizophrenia-like psychosis in a clear sensorium. The defining combination, as set out in Kaufman's Clinical Neurology for Psychiatrists, is a chronic psychotic illness that develops interictally, in clear consciousness, and that phenomenologically resembles schizophrenia. Each of those three qualifiers is doing work. *Chronic* distinguishes it from the acute, self-terminating epilepsy psychoses and signals that this is an illness measured in months and years rather than hours and days. *Interictal* places it off the seizure timeline: the psychosis is not the aura, not the fit and not the immediate aftermath, but a state that persists in the gaps between seizures and does not track them. *Clear consciousness* is the anchor that keeps it apart from every clouded, confused or deliriant peri-ictal picture, because the patient is alert, oriented and coherent in sensorium even while floridly deluded. The resemblance to schizophrenia is genuine and not merely superficial, which is exactly why the discrimination from the primary disorder had to become a discipline of its own. Framing it as a secondary psychosis, real and schizophrenia-like but arising on an epileptic brain, is what justifies taking a seizure history in every apparently primary psychosis and explains the temporal signature that no primary schizophrenia carries. Current classifications place the epilepsy psychoses among the secondary presentations rather than treating them as schizophrenia proper, which keeps the clinician's attention on the underlying epilepsy as well as on the psychotic state in front of them.

17.2 How common it is, and the epilepsy-schizophrenia association

Uncommon in absolute terms, but far above the population rate. Psychosis is not the commonest psychiatric complication of epilepsy, sitting well below the mood and anxiety disorders in frequency, yet it is disproportionately important because it is severe, chronic and closely tied to the epilepsy itself. The Maudsley Prescribing Guidelines give the point prevalence of psychosis among people with epilepsy as roughly 5%, an order of magnitude above the general population, with depression and anxiety more common still. The substrate matters as much as the headline figure. Temporal lobe epilepsy, whose anatomy and semiology are covered in their own right in this chapter, is the principal setting, and Kaplan and Sadock's Comprehensive Textbook of Psychiatry reports that up to 9% of adults with temporal lobe epilepsy develop an associated interictal psychosis. There is a laterality signal as well: the psychosis is described as more common with a left, dominant-hemisphere temporal focus than with a right-sided one, a finding that has fed the long-running argument that dominant limbic pathology is particularly psychotogenic. Stepping back from epilepsy to the population level, Lishman's Organic Psychiatry cites the Danish register work of Qin and colleagues, in which a history of epilepsy carried a relative risk of schizophrenia of 2.48 (95% CI 2.20 to 2.80), a two- to three-fold elevation. That register finding is one strand of the broader bidirectional relationship between epilepsy and psychiatric disorder, which is set out in full in its own place in this chapter; what it

establishes for the present topic is that the association is real and directional, not an artefact of shared ascertainment.

17.3 Age of onset and the temporal signature

Epilepsy in childhood, psychosis in the fourth decade, a long gap between. The most examinable feature of this illness is not a symptom at all but a chronology. Kaufman's Clinical Neurology for Psychiatrists places the typical onset of the interictal psychosis at **30 to 40 years**, in patients whose epilepsy began in childhood, characteristically in the first decade of life, so that the psychosis emerges decades after the seizures first appeared. That gap is not incidental. Lishman's Organic Psychiatry frames it as an interval of roughly 10 to 15 years separating the onset of epilepsy from the onset of the schizophrenia-like psychosis, a latency confirmed by report after report since the original description. The classic series is that of Slater, Beard and Glithero (1963), whose sixty-nine patients showed a mean age at onset of the psychosis of 30 years after a mean epilepsy duration of 14 years. The practical force of the chronology is that a first psychosis in a young or middle-aged adult who has carried epilepsy since childhood should raise this diagnosis before a first break of primary schizophrenia is assumed. The accompanying table sets the temporal anchors side by side.

TEMPORAL ANCHOR	REPORTED FIGURE	SOURCE
Age at onset of the psychosis	typically 30 to 40 years, epilepsy from childhood	Kaufman's Clinical Neurology
Latency, epilepsy onset to psychosis onset	about 10 to 15 years	Lishman's Organic Psychiatry
Classic series (Slater, Beard and Glithero)	mean onset 30 years, after 14 years of epilepsy	Lishman's Organic Psychiatry

The temporal signature from three sources: childhood-onset epilepsy, a latency of roughly a decade and a half, and a psychosis declaring itself around the fourth decade. The chronology is the discriminator that a single mental state examination cannot supply.

17.4 Clinical picture: a predominantly positive psychosis

Loud positive symptoms in a clear head. Kaufman's Clinical Neurology for Psychiatrists describes a picture that is predominantly made of **positive symptoms**: **persistent hallucinations**, **paranoid and persecutory delusions**, and a tendency to **social isolation**, all occurring in otherwise clear consciousness. The hallucinations are sustained rather than fleeting, the delusional content is characteristically persecutory, and the withdrawal from social contact can be marked, so that the cross-sectional impression genuinely reads as schizophrenia. This is the crux of the topic and the source of most of its difficulty: nothing in the roster of positive symptoms reliably points away from a primary psychosis, and a cardinal schizophrenic symptom in such a patient does not exclude epilepsy. What differs lies in the longitudinal frame rather than the cross-section, which is precisely why the systematic discrimination of this entity from primary schizophrenia is developed in its own dedicated topic and is not reopened here, and why the phenomenology of primary schizophrenia belongs to the Schizophrenia: aetiology, phenomenology, classification and course notes. For the purposes of this entity, the message is that the interictal psychosis presents as a persistent, hallucinated, paranoid state in a patient who

is fully alert, and that the clarity of the sensorium, far from arguing against psychosis, is one of its defining features.

MNEMONIC

Three P's in a clear head

The interictal-psychosis picture in four beats. **P**ersistent hallucinations. **P**aranoid, persecutory delusions. **P**oor social contact, with withdrawal into isolation. All of it in a **clear** sensorium, interictally and chronically. Each element is a feature the source actually states; the clear head is the beat that keeps it off the peri-ictal timeline.

TRAP **Expecting a clouded sensorium.** Candidates lose the mark by describing this psychosis as arising in confusion or drowsiness, importing the clouding of the ictal and postictal states into an illness defined by its absence. By definition the interictal psychosis occurs in clear consciousness; a clouded patient is having a different syndrome. State the clarity of the sensorium explicitly, because it is examined as a discriminator, not as an afterthought.

17.5 Who is at risk: severity of epilepsy and brain damage

The risk markers read as an index of epilepsy severity. The associations that predict interictal psychosis are, with striking consistency, markers of how severe and how damaging the underlying epilepsy has been, which is the single most useful way to remember them. Kaufman's Clinical Neurology for Psychiatrists groups them as follows.

- Childhood onset of the epilepsy, the same early start that opens the long latency to psychosis.
- A demonstrable neurological abnormality or evidence of brain damage.
- Low intelligence.
- Frequent seizures and multiple seizure types.
- Treatment with multiple antiepileptic drugs, itself a marker of refractory disease rather than a cause.
- A prior postictal psychosis.

Reading these as a single construct rather than a checklist is what earns the mark: they describe a brain seized early, often, in more than one way, and with structural or cognitive compromise, and it is that cumulative burden that raises the psychotic risk. The point about multiple antiepileptic drugs deserves emphasis because it is easily misread. Polytherapy here is a proxy for hard-to-control epilepsy, not evidence that the drugs are psychotogenic; the drug-attributable and forced-normalisation routes to psychosis are separate constructs taught elsewhere. The last item on the list draws an explicit bridge between the acute and chronic psychoses of epilepsy. Lishman's Organic Psychiatry, citing the work of Tarulli and colleagues, records that about **15%** of patients with a postictal psychosis, the brief seizure-locked form taught as its own entity, eventually go on to develop a chronic interictal psychosis. That progression argues that seizure-related mechanisms drive the chronic illness, and clinically it makes a recurrent postictal psychosis a longitudinal warning, not only an acute problem.

17.6 Course and prognosis

More often chronic than not, with onset speed the best guide. The name of the entity already tells the story of its course, but the figures put a shape on it. Lishman's Organic Psychiatry records that nearly half of Slater's original series ran a chronic course, and that in a longer follow-up Onuma found **64% remaining chronic over ten years**. Once established, this psychosis tends to persist rather than resolve, which is a large part of what separates it prognostically from the brief peri-ictal psychoses that recover fully. The most useful single prognostic pointer is the tempo of onset: an **acute onset carries a better prognosis** than an insidious one, so the manner in which the illness begins is worth documenting from the first contact. The chronicity has direct consequences for how the patient is followed. An illness that persists for years in a person who already carries chronic epilepsy is a long-term management problem in two domains at once, the seizures and the psychosis, and neither can be allowed to drift while the other is attended to. It also reframes the goal of treatment: the realistic aim in the majority is sustained control and preserved function over years, not a brief course of medication with an expected cure. The minority whose illness is acute and self-limiting are the exception, and identifying them early avoids over-committing them to lifelong treatment.

17.7 Management

Two illnesses in one patient, treated in parallel. Management rests on treating the psychosis while never taking the eye off the epilepsy, and it is guided by the ILAE and NICE epilepsy guidelines for seizure control, by the Maudsley Prescribing Guidelines for the antipsychotic itself, and by the epilepsy-specific psychiatric-management literature that bridges the two. The **LOCUS** of care is, for a chronic illness, predominantly the outpatient clinic, with inpatient admission reserved for acute exacerbation, behavioural risk or the acutely disturbed presentation that reaches an emergency setting. The **focus** shifts by phase. The acute target is the florid positive symptoms and any risk they generate; the short-to-mid-term target is a durable antipsychotic response together with optimisation of seizure control; and the long-term target, given the chronic course, is maintenance of remission, relapse prevention and the rehabilitation of a patient whose social contact may have narrowed. The **modus** spans every mode of treatment. Pharmacologically, Kaplan and Sadock's Comprehensive Textbook of Psychiatry notes that the majority of patients require antipsychotic treatment, although up to 15% remit without any, and that there is no clear efficacy difference between first- and second-generation agents, so the choice turns on tolerability and, crucially, on effect on the seizure threshold. That specific decision, which antipsychotic to choose against its propensity to lower the seizure threshold, is handled in its own dedicated topics and is not reopened here. Beyond the drug sit the psychological, social and procedural modes: engagement and psychoeducation, attention to the social isolation that is part of the illness, family involvement, and above all the tightening of antiepileptic treatment as a lever that acts on the shared seizure substrate.

PITFALL

Because a slice of these patients recover without medication, it is tempting to watch and wait in an alert, non-dangerous patient and defer the antipsychotic. The numbers do not support it: only a small minority remit untreated, and the great majority need an antipsychotic to get better. Waiting turns a treatable chronic psychosis into a longer one and lets the social withdrawal entrench. Treat the psychosis; do not gamble on the exception.

GOOD TO REMEMBER

When a person with long-standing, childhood-onset epilepsy develops a persistent paranoid-hallucinatory psychosis in a clear sensorium, do not reflexively file it as a first episode of primary schizophrenia. Start an antipsychotic, chosen with the seizure threshold in mind, and at the same time review and optimise the antiepileptic regimen. You are managing the psychosis and the epilepsy together, and the seizure control is part of the psychiatric treatment, not a separate errand.

IN INDIA

Two features of Indian practice sharpen this problem. The epilepsy treatment gap is wide, so childhood-onset epilepsy is often poorly controlled for years, which is precisely the severity profile that breeds the chronic interictal psychosis; optimising long-neglected seizure control is frequently the single most valuable intervention. Second, longitudinal records and prior electroencephalography are often unavailable, so a chronic schizophrenia-like picture is easily mislabelled as primary schizophrenia when no reliable seizure history is to hand. The practical safeguard is to take an active seizure history in every chronic psychosis rather than waiting for one to be volunteered, and to weigh the double stigma of epilepsy and mental illness when planning long-term, community-based follow-up.

17.8 Where it sits among the psychoses of epilepsy

The chronic, seizure-independent member of the family. Keeping this entity distinct from its neighbours is best done, once again, on the timing axis that organises the whole group. It is not the ictal psychosis, which is itself the electrographic seizure and has no interval at all. It is not the postictal psychosis, which follows a cluster of fits across a lucid interval and then remits, though a minority of those cases progress into this one. It is not the alternative psychosis of forced normalisation, a distinct construct in which the disturbance appears as the electroencephalogram normalises and seizures cease. And where an interictal presentation is dysphoric and affective rather than psychotic, the interictal dysphoric disorder of epilepsy is the neighbour to consider. The formal business of separating any of these from a primary schizophrenic illness is developed in its own dedicated topic, and the primary illness itself belongs to the Schizophrenia: aetiology, phenomenology, classification and course notes. Hold the entity by its signature and it stops blurring into the family around it.

RULE **Chronic psychosis, clear sensorium, running on between seizures in long-standing childhood epilepsy, is the schizophrenia-like psychosis of epilepsy.** Every word of that line is load-bearing: chronicity separates it from the acute peri-ictal psychoses, the clear sensorium separates it from the clouded ones, the interictal, seizure-independent course separates it from the ictal and postictal forms, and the long latency from childhood-onset epilepsy is the chronology no primary schizophrenia supplies.

up to 9%	about 5%	10 to 15 yr	2.48
INTERICTAL PSYCHOSIS IN TEMPORAL LOBE EPILEPSY	POINT PREVALENCE OF PSYCHOSIS IN EPILEPSY	EPILEPSY-TO-PSYCHOSIS LATENCY	RELATIVE RISK OF SCHIZOPHRENIA GIVEN EPILEPSY

The schizophrenia-like (chronic interictal) psychosis of epilepsy is a chronic, positive-symptom psychosis in clear consciousness, arising a decade or more after childhood-onset epilepsy, most often in temporal lobe epilepsy; most patients need an antipsychotic, the course is usually chronic, and the epilepsy must be treated alongside the psychosis.

18 · Forced normalisation (Landolt phenomenon) and alternative psychosis

This is the one epilepsy psychosis that punishes success. In almost every other peri-ictal and interictal syndrome the psychiatric disturbance tracks the seizures, so that controlling the fits tends to calm the behaviour. **Forced normalisation** inverts that comfortable rule. It names the emergence of psychotic, or less often severe affective, symptoms precisely as the seizures come under control and the electroencephalogram quietens, so that the very sign the neurologist has been working towards, a clean record, can arrive hand in hand with a new psychiatric illness. The eponymous label, the **Landolt phenomenon**, and its clinical twin, **alternative psychosis**, describe the same counter-intuitive reciprocity from two different vantage points. The marks are here because the idea is conceptually clean, historically rich, and a favourite of examiners probing whether a candidate can hold that seizures and psychosis might, in a subset of patients, behave as antagonists rather than allies.

The concept is old and its status contested, but its practical importance is out of proportion to how rarely it is convincingly documented. If it is real, the pursuit of complete seizure freedom is not an unqualified good in every patient, and the instruction that falls out of it is safe whatever one makes of the mechanism: watch the patient whose seizures have just stopped.

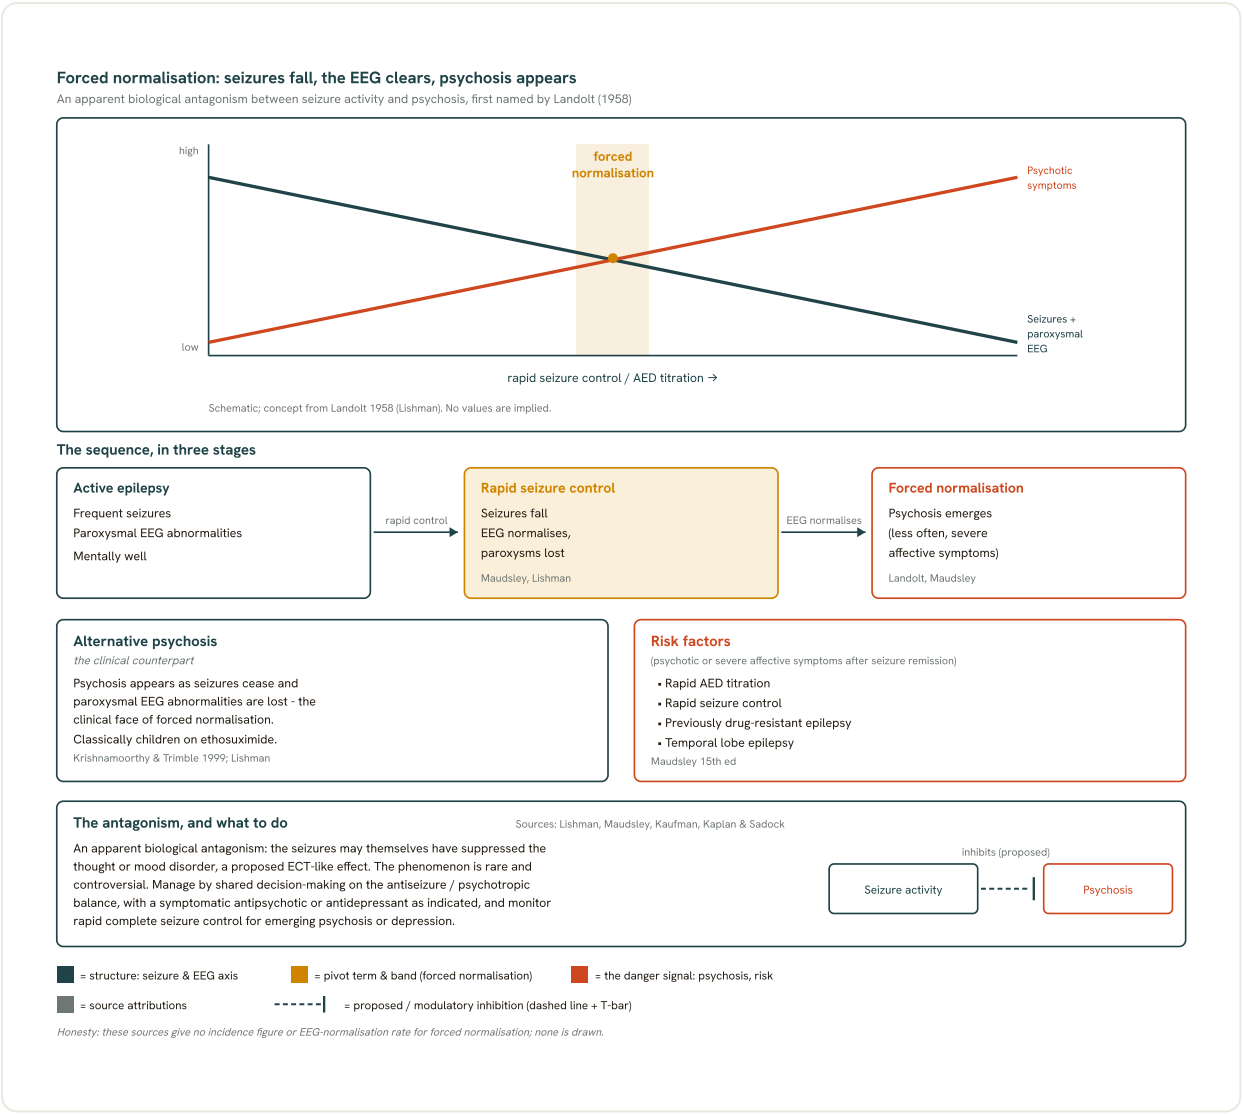


Fig 18.4 Forced normalisation and alternative psychosis. The EEG normalises and paroxysmal abnormalities are lost just as psychotic (or, less often, severe affective) symptoms appear, an apparent biological antagonism between seizures and psychosis first named by Landolt in 1958; alternative psychosis is its clinical counterpart, classically in children whose seizures cease on ethosuximide.

18.1 What forced normalisation is

Psychiatric symptoms that rise as seizures fall. The core definition is a temporal inversion. Forced normalisation is the appearance of psychosis, or a severe affective disturbance, at the moment when seizure frequency drops and the EEG loses its epileptiform activity, so that the improvement in the electrical disorder and the arrival of the mental disorder are yoked together in time. The Maudsley Prescribing Guidelines place it among the para-ictal episodes for exactly this reason: it is defined not by any particular content of the psychosis but by its relationship to a falling seizure burden. The essential clause, and the one that separates it from every look-alike, is that the seizure remission comes first and the psychiatric episode follows it. A psychosis that runs alongside continuing frequent seizures is not this entity; a psychosis that appears only once the fits have been suppressed is its candidate. Reading the definition this way keeps the attention on sequence rather than on symptom. The picture may be paranoid, grandiose, affectively coloured or frankly depressive, and the phenomenology alone will not name the syndrome for you; what names it is that the disturbance emerged out of a quietening brain rather than a convulsing one.

It is a syndrome of timing and reciprocity, not of a fixed clinical face, and its examinable heart is the paradox that the patient got psychiatrically worse as they got neurologically better.

18.2 Landolt and the origin of the term

An observation born at the EEG machine. The term is a piece of electroencephalographic history, and its provenance is examinable in its own right. It was **Landolt** who coined it, in **1958**, from serial EEG recordings taken through the psychotic episodes of patients with epilepsy, as documented in Lishman's Organic Psychiatry. The critical word is normalisation, and it is meant literally: Landolt was describing the record, not the patient. What he saw was that in certain individuals the EEG became normal, its paroxysmal features disappearing, at the very time the psychosis declared itself, and returned to its abnormal epileptiform pattern as the psychosis lifted. The adjective forced captures the sense that this normalisation was not the ordinary quiet of a well person but a suppression that seemed to buy the clean trace at the cost of the mind. Holding the origin in mind protects against a common misreading: forced normalisation is not a clinical description that happens to mention an EEG but a fundamentally EEG-defined phenomenon, an observation about what serial records do during psychotic episodes in epileptic patients. That distinction matters because the entity was later re-approached from the bedside, where the electroencephalogram is not always the tool in hand, and the two descriptions have to be kept aligned.

18.3 Alternative psychosis: the clinical counterpart

The same reciprocity, seen from the bedside. Because the neurophysiologist's EEG-defined term does not travel easily into everyday clinical language, the phenomenon acquired a clinical counterpart. Alternative psychosis is the name for the same reciprocal event described in behavioural terms: the psychosis appears as overt seizures cease, so that seizures and psychosis seem to alternate, one giving way to the other. Where a record happens to be taken the paroxysmal EEG abnormalities may normalise in step, which is the forced-normalisation counterpart, but the clinical term is defined by the alternation of seizures and psychosis and does not require documented EEG normalisation, which is exactly what lets it be used at the bedside where no serial EEG is in hand. Lishman's Organic Psychiatry attributes the modern framing to Krishnamoorthy and Trimble in **1999**, who set the clinical concept alongside Landolt's electrographic one, and its classic illustration is a childhood one: children whose absence seizures are rendered fully controlled on **ethosuximide** and who then develop a psychiatric disturbance as the fits vanish. The word alternative is doing careful work: it signals substitution, a swap of one state for another rather than mere coincidence, and it is the behavioural expression of what Landolt read off the trace. Keeping the pair straight is a small examination skill in itself, and the accompanying figure sets the two descriptions side by side.

AXIS OF DESCRIPTION	FORCED NORMALISATION (LANDOLT)	ALTERNATIVE PSYCHOSIS
Anchored in	the electroencephalogram	the clinical seizure state
The change that defines it	paroxysmal EEG abnormality disappears	overt seizures cease
Vantage point	the neurophysiologist reading serial records	the clinician watching the patient
Classic illustration	chronic, often drug-resistant epilepsy in adults	a child rendered seizure-free
What it asserts	one reciprocal phenomenon, seizures and psychosis alternating, described two ways	

Forced normalisation and alternative psychosis are the electrographic and the clinical descriptions of a single reciprocity; the term you reach for depends on whether the EEG or the seizure state is the thing in front of you.

18.4 The proposed mechanism: a biological antagonism

Seizures and psychosis as opposing states. The organising hypothesis is that, in some individuals, seizure activity and psychiatric disturbance stand in a relationship of **biological antagonism**, so that active epileptiform activity somehow holds the thought or mood disorder in check and its removal releases it. Lishman's Organic Psychiatry frames this as the working mechanism behind both the electrographic and the clinical descriptions: the seizures may have been suppressing an underlying psychiatric disorder, and abolishing them lifts that suppression. The most cited version of the idea invokes a parallel with electroconvulsive therapy. A seizure is, at the level of the brain, a massive synchronous discharge with recognised antidepressant and antipsychotic effects when delivered therapeutically, and the antagonism hypothesis proposes that a patient's own spontaneous seizures might exert something of the same stabilising influence, so that a brain deprived of them loses an endogenous treatment. This is a mechanistic proposal rather than an established pathway, but it has real explanatory reach. It accounts for the timing, since the disturbance follows the loss of seizures; it accounts for the affective as well as the psychotic presentations, since the postulated protective effect is not specific to one symptom class; and it fits the observation that the phenomenon is seen when control is achieved rapidly, the putative endogenous stabiliser withdrawn abruptly rather than tapered. Its value is that it converts a list of disconnected facts into a single story, and it is the point examiners probe when they ask why controlling seizures might, paradoxically, be the thing that made the patient ill.

RULE **Better seizure control can buy worse mental health.** In a susceptible minority, psychosis or a severe affective disturbance appears as seizures remit and the EEG clears, which is the reverse of the usual expectation that controlling fits calms behaviour. The corollary is practical rather than nihilistic: rapid, complete seizure control is a setting to watch, not a result to celebrate uncritically.

18.5 Who is at risk

A recognisable clinical setting, not a random event. Although uncommon, the phenomenon does not strike at random, and the reported cases share a describable profile that lets a clinician

anticipate it. The Maudsley Prescribing Guidelines set out the features that recur when psychotic or severe affective symptoms follow seizure remission, and they cohere into a picture of a hard-to-control epilepsy just brought sharply to heel. Four features travel together as the risk profile for a forced-normalisation reaction:

- Rapid up-titration of the antiseizure drug, so that the seizure burden falls steeply rather than gradually.
- Rapid and complete control of seizures, the abrupt achievement of freedom from fits.
- An epilepsy that had previously been medication-resistant, so that the sudden control is itself unusual.
- Temporal lobe epilepsy, whose anatomy, semiology and psychiatric salience are taught in their own right elsewhere in this chapter.

Each element makes sense through the antagonism lens. If spontaneous seizures were supplying an endogenous stabilising effect, the danger is greatest where that effect is removed quickly and wholly, which is what rapid titration and rapid complete control describe. The prior drug resistance marks the patients in whom sudden, thorough control is both possible and startling to the system, and the temporal lobe association ties the phenomenon to the limbic structures through which seizure activity most readily reaches mood and belief. The risk factors are therefore less a list to recite than a situation to recognise: a previously refractory patient in whom a briskly escalated regimen has just abolished the fits.

MNEMONIC

Three R's and a T

The setting in which a forced-normalisation reaction tends to be reported. **R**apid antiseizure-drug titration, **R**apid and complete seizure control, a **R**efractory (previously drug-resistant) epilepsy, and **T**emporal lobe epilepsy. Each letter is a risk feature the topic actually names, not an invented extra.

18.6 A rare and genuinely controversial entity

Real enough to teach, contested enough to qualify. The status of the phenomenon is itself examinable, since the safest answer neither dismisses it nor overstates it. Kaplan and Sadock's Comprehensive Textbook of Psychiatry describes forced normalisation as a rare and controversial phenomenon, in which psychosis coincides with improved seizure control achieved by any route: medication, vagus nerve stimulation, or epilepsy surgery. That framing carries two useful teaching points. The first is the breadth of trigger: it is not a drug effect in the narrow sense, since the psychosis can follow control secured by a stimulator or by resection just as readily as by tablets, which is what one would expect if the operative variable is the seizure suppression rather than the means used to achieve it. The psychiatric outcomes of epilepsy surgery are a topic in their own right elsewhere in this chapter; this is only their point of contact with forced normalisation. The second point is the controversy, and it deserves to be understood rather than merely stated. The phenomenon is hard to establish because its classic definition depends on serial EEG capture across an episode that is itself uncommon; ascertainment is biased toward the

cases that happened to be recorded, causality is difficult to separate from coincidence in a population already prone to psychosis, and the literature leans on case reports. None of this makes it unreal, and its corollary of vigilance stands regardless, but it does mean forced normalisation should be presented as an established concept of uncertain frequency and contested mechanism rather than as a settled quantity.

1958	1999	rare
LANDOLT NAMES THE PHENOMENON	ALTERNATIVE PSYCHOSIS REFRAMED CLINICALLY	AND GENUINELY CONTROVERSIAL IN STATUS

18.7 Recognition and the differential

Made on sequence, defended against its mimics. Because there is rarely a serial EEG to hand, recognition in practice is a clinical act that rests on reconstructing the sequence: an epilepsy, often refractory, that was brought under control, and a psychiatric episode that appeared as the control was achieved. Kaufman's Clinical Neurology for Psychiatrists turns this into an actionable caveat, advising that clinicians monitor for emerging psychosis or depression the patient who rapidly achieves complete seizure control after a change of antiseizure regimen that suppresses the seizures and the EEG. The differential is where the marks and the errors lie. The disturbance must be separated from a direct psychiatric side effect of the antiseizure drug that was escalated, the psychiatric side effects of individual antiepileptic drugs being a distinct topic in their own right; the two can look identical at the bedside and are told apart by whether the operative change was the drug itself or the remission it produced. It must be separated from the other psychoses of epilepsy that arrive on a different schedule: the postictal psychosis that erupts after a lucid interval following a cluster of fits, and the schizophrenia-like psychosis of epilepsy, the chronic interictal picture that runs on independently of seizure control and is taught as its own entity elsewhere in this chapter. Where the presentation is affective and interictal, the interictal dysphoric disorder of epilepsy is the neighbour to consider. The systematic business of distinguishing any epileptic psychosis from primary schizophrenia, on preserved affect, premorbid function and the relation to seizures, is developed in its own dedicated discussion, and the phenomenology of the primary illness itself belongs to the Schizophrenia: aetiology, phenomenology, classification and course notes.

TRAP **Blaming the drug instead of the remission.** When a patient turns psychotic soon after an antiseizure drug is pushed up and the fits stop, the reflex is to record a psychiatric side effect of that drug and switch it. Forced normalisation attributes the disturbance not to a direct pharmacological action but to the seizure suppression itself, and the distinction changes whether one would ever deliberately accept less complete control. That is the discriminating point examiners are listening for.

PITFALL

The opposite error is over-diagnosis. Not every person with epilepsy who becomes psychotic after starting treatment has forced normalisation; the phenomenon is uncommon, and reaching for it whenever seizures improve and a mental disturbance appears risks missing a coincidental primary psychosis, a direct antiepileptic side effect, or a postictal psychosis with a quite different time course. Reserve the label for the genuine temporal coupling of remission and a quietening EEG with the new symptoms.

GOOD TO REMEMBER

Whenever a change of regimen abolishes seizures quickly, especially in someone whose epilepsy had been hard to control, keep the mental state under review over the following weeks and ask the family to report emerging suspiciousness, elation or low mood. It is the low-cost instruction the topic exists to teach, and it catches the problem early without committing anyone to a change in the epilepsy plan before it is needed.

IN INDIA

Read the lesson forwards, not backwards. With a large epilepsy treatment gap, the dominant problem across most of Indian practice is under-treatment, so forced normalisation is never a reason to aim for anything less than good seizure control. Where it earns attention is the tertiary-centre patient whose refractory epilepsy is suddenly brought to heel by an optimised regimen, vagus nerve stimulation or surgery; because serial EEG follow-up is often not feasible, recognition then rests on clinical vigilance and a well-briefed family rather than on the record.

18.8 Management

The management question is a balance, not a protocol. What makes this topic unusual is that the treatment and the possible cause are the same lever, a genuine trade-off rather than an algorithm. The organising framework rests on seizure-control standards from the ILAE and the NICE epilepsy guidelines for the epilepsy itself, and on the Maudsley Prescribing Guidelines together with the wider epilepsy-psychiatry literature for the psychiatric episode. The guiding principle those sources converge on is shared decision-making with the patient and their carers about how to proceed with the antiseizure and psychotropic drugs, precisely because the values at stake, freedom from seizures against freedom from psychosis, belong to the patient to weigh.

The LOCUS of care follows severity. Most patients can be held and monitored as outpatients; a florid or high-risk psychotic or affective episode may warrant inpatient admission, and acute behavioural disturbance may present first to an emergency setting. The focus shifts by phase. The acute task is to contain the psychiatric symptoms and secure safety; the short-to-mid-term task is to settle the antiseizure-versus-psychotropic balance deliberately rather than reflexively; and the long-term aim is a sustainable equilibrium in which seizures and mental state are both tolerable. The modus spans every mode of treatment. On the drug side, symptomatic treatment with an antipsychotic or an antidepressant is used as the presentation indicates, while the separate question of which antipsychotic to choose against its effect on the seizure threshold, and the

mechanism of that threshold-lowering, are handled in their own dedicated topics and are not reopened here. The antiseizure decision itself is the pivot, and it is where the shared decision belongs: whether to hold the regimen and treat the psychiatric episode symptomatically, or, uncommonly and only by agreement, to accept marginally less complete seizure control, is a judgement made with the patient and family rather than imposed. Around the drugs sit the psychological and social modes, chiefly education of the patient and carers about the paradox they are living through, so that neither the improvement in seizures nor the emergence of symptoms is misread.

> Forced normalisation, the Landolt phenomenon, is the emergence of psychosis or severe mood disturbance as seizures are brought under control and the EEG normalises; its bedside counterpart is alternative psychosis, the working idea is a reciprocal antagonism between seizures and psychiatric illness, and the safe clinical rule is to watch the mental state of any patient whose fits have just been abolished.

19 · Differentiating epileptic psychosis from primary schizophrenia

Why the marks are here. The examiner who asks you to separate the psychosis of epilepsy from primary schizophrenia is not testing whether you can recognise psychosis, because both patients are psychotic; the test is whether you understand that a psychotic cross-section is a starting point and not a diagnosis. The problem is genuine. The chronic psychosis that arises between seizures in long-standing epilepsy, taught in full as the schizophrenia-like psychosis of epilepsy, can reproduce the surface of schizophrenia so faithfully that a single mental state examination will not tell the two apart. What separates them is almost everything the mental state examination does not capture in one sitting: the patient's affect and personality over time, the trajectory of the illness, the family background, and above all the relationship of the psychosis to a history of seizures. This topic owns that discrimination. It does not re-teach the primary illness, whose phenomenology, aetiology and course belong to the Schizophrenia: aetiology, phenomenology, classification and course notes, and it does not re-describe the epileptic entity itself, which is set out where the schizophrenia-like psychosis of epilepsy is taught.

Two boundaries keep the topic clean. The discrimination here concerns the chronic interictal presentation, the psychosis that sits between seizures and persists. The brief, seizure-locked forms are a different problem: the ictal psychosis that is itself the electrographic seizure, and the postictal psychosis that follows a cluster of fits after a lucid interval, are separated from primary schizophrenia chiefly by their timing rather than by the features discussed below, and each is dealt with as its own entity, as is the alternative psychosis of forced normalisation. So when this chapter speaks of differentiating epileptic psychosis from schizophrenia, it means the chronic, interictal, schizophrenia-like picture, because that is the one that genuinely imitates the primary disorder and forces the question.

19.1 The resemblance is real: symptoms alone will not separate them

The two overlap almost completely at the level of individual symptoms. The canonical description comes from **Slater** and colleagues, whose series of **69 patients** with epilepsy who had developed a schizophrenia-like illness established that these patients showed all the cardinal symptoms of schizophrenia, delusions, hallucinations and passivity phenomena among them, and yet in atypical combinations. In other words, nothing in the roster of positive symptoms was missing, and a cardinal schizophrenic symptom in such a patient does not point away from epilepsy. What was atypical was the pattern: the way the symptoms hung together and the relative sparing of the features that dominate severe primary schizophrenia. This is the single most important fact to carry into the examination, because the instinct of the trainee is to hunt for a symptom that rules schizophrenia in, and there is no such symptom to be found.

Understanding why the resemblance is so close also tells you where to look for the difference. Both illnesses generate psychosis, but they arrive at it by different routes. Primary schizophrenia is understood as a pervasive disorder of the developing brain that erodes the personality and the capacity for affective and social engagement, often before the first frank psychotic symptom appears. The psychosis of epilepsy is a symptomatic psychosis: a productive disturbance grafted onto a brain and a personality that were, until the illness declared itself, developing normally, and driven by a focal epileptic process rather than by the diffuse liability that underlies the primary disorder. Because the substrate differs, the surface can be identical while the depths are not, and it is in those depths, in affect, personality, course, heredity and the temporal link to seizures, that the discrimination is actually made.

TRAP **Believing that a fully schizophrenic cross-section excludes an epileptic cause.** Candidates reason that florid, well-systematised delusions or vivid auditory hallucinations must indicate the primary disorder. Slater's patients had exactly these features. The presence of cardinal schizophrenic symptoms is entirely compatible with the psychosis of epilepsy; it is their atypical combination and the surrounding history, not their mere presence, that carry the diagnostic weight.

19.2 Preserved affect and easy rapport at the bedside

Warm affect and interpersonal openness are the most useful bedside signs. The texture of affect and relatedness is where the difference first shows itself in a single interview. In the psychosis of epilepsy the affect remains, in relative terms, warm and responsive. The **preserved affect** of these patients is one of the oldest and most reliable points of separation, because the loss of affective response that characterises primary schizophrenia is here neither as early nor as marked. The patient continues to show emotional colouring appropriate to the content of thought and keeps the capacity to make affective contact with the examiner, rather than presenting the pervasive blunting and incongruity that dissolve rapport in the primary illness.

This warmth extends into the interpersonal domain. Clinically these patients tend to be approachable and willing rather than guarded, more cooperative and notably **less suspicious** than patients with primary schizophrenia. Where the schizophrenic patient is often withdrawn and hard to engage, the patient with an epileptic psychosis is frequently friendly and

forthcoming even while holding florid delusional beliefs. This is not a soft or impressionistic sign. It reflects the underlying difference between the two disorders, namely that the epileptic psychosis leaves the affective and relational core of the personality comparatively intact while superimposing a productive psychosis on top of it. The practical corollary is direct: a psychotic patient who remains emotionally engaged and cooperative, whose delusions do not come wrapped in the deep suspiciousness and affective distance of the primary illness, is a patient in whom you should be asking about seizures rather than closing the diagnosis.

19.3 A better premorbid baseline and a course that does not deteriorate

Two temporal features, one looking back and one looking forward. The impression the affect gives is reinforced from both directions in time. Looking backward, the **premorbid personality** is relatively well preserved: premorbid function is generally better than in primary schizophrenia, and a prominent personality abnormality in the years before the psychosis declared itself is uncommon. The insidious premorbid slide so often seen before a first episode of the primary illness, the quiet withdrawal and the erosion of drive and social function, is characteristically absent, and by history the patient was functioning reasonably well as a person before the epilepsy-related psychosis emerged.

Looking forward, the course diverges even more clearly. Patients with the epileptic psychosis do not show the **progressive deterioration** that is characteristic of schizophrenia. Over years the primary illness tends towards an accumulating defect state, with deepening negative symptoms and functional decline; the epileptic psychosis carries no such autonomous downhill trajectory. Whatever decline the patient shows tends to track the epilepsy and its consequences rather than an independent psychotic process. This is why the longitudinal view is so powerful. Give the two conditions time and the primary disorder tends to declare itself through deterioration, while the epileptic psychosis holds a comparatively steady line.

> **GOOD TO REMEMBER**
>
> The preserved personality and the non-deteriorating course are not merely diagnostic curiosities; they carry a better functional prognosis than primary schizophrenia and they change what you tell the family. Counselling that assumes the relentless trajectory of the primary illness misrepresents this patient's likely course, and a pessimism imported wholesale from schizophrenia is neither accurate nor kind.

19.4 The family history does not carry a schizophrenic loading

The pedigree speaks to aetiology, not presentation. Heredity gives one of the cleaner points of separation precisely because it bypasses the symptom picture altogether. Patients with the interictal psychosis of epilepsy do not show the increased **familial incidence of schizophrenia** seen in the primary disorder. Primary schizophrenia carries a well-recognised familial loading, with an increased occurrence of the illness among relatives; the epileptic psychosis does not sit on that genetic background. A family history of psychosis, where it is present, behaves at most as a modest risk factor for developing psychosis in the setting of epilepsy, not as the schizophrenia-level loading that the primary illness implies.

The reasoning follows once the aetiologies are separated. If the psychosis is being driven by an acquired, structural, epileptic process, then the inherited liability that runs in the families of people with primary schizophrenia should not be enriched in the families of these patients, and it is not. For the clinician this has two uses. It is a discriminator, in that a strong family history of schizophrenia shifts the balance towards the primary diagnosis; and it is a counselling point, because the heritable risk you would quote to the relatives of a person with primary schizophrenia does not transfer to the relatives of a patient whose psychosis is a complication of epilepsy.

19.5 The psychosis is grafted onto years of established epilepsy

Epilepsy first, by years, and then the psychosis. The most powerful contextual discriminator, and the one an examiner will reward most, is the temporal relationship between the psychosis and the seizure disorder. The interictal psychosis of epilepsy characteristically emerges only after the epilepsy has been established for a long time, on the order of a mean interval of about **14 years** of active seizures before the psychosis appears. The epilepsy comes first, by years, and the psychosis is grafted onto it. Primary schizophrenia has no such antecedent; there is neither a requirement for, nor usually any history of, a preceding seizure disorder of long standing.

This long latency matters both diagnostically and conceptually. Diagnostically, it means the single most discriminating question you can ask is about the sequence: was there a well-established epilepsy, active over many years, before the psychosis began? An affirmative answer, in a patient whose affect and personality are relatively preserved and whose course is not deteriorating, points firmly towards the epileptic psychosis. Conceptually, the long interval suggests a slow, secondary process, the gradual consequence of years of recurrent seizure activity, rather than the developmental disorder that underlies the primary illness. The mechanism linking chronic seizure activity to behavioural disturbance belongs to the chapter's account of the neurobiology of epilepsy and behaviour and is only named here, not re-derived.

> **PITFALL**
>
> The commonest way this diagnosis is missed is a seizure history that was never taken. A patient presenting with a first psychotic illness is placed on a primary-psychosis pathway, and an established epilepsy, sometimes nocturnal, sometimes long-standing focal seizures the patient no longer thinks to mention, is never elicited. Ask every psychotic patient directly about seizures, blackouts and their treatment. The temporal discriminator is worthless if the epilepsy is never uncovered in the first place.

19.6 Phenomenological texture: thought disorder and catatonia

Positive symptoms dominate, without the disorganisation and the motor phenomena. Although symptomatology cannot separate the two disorders on the presence or absence of cardinal symptoms, the emphasis within the symptom picture does differ, and two negatives are worth knowing. First, the interictal psychosis often shows a relative absence of **formal thought disorder**. The productive positive symptoms, the delusions and the hallucinations, tend to dominate, while the disorganisation of thought and language that can be so prominent in

primary schizophrenia is comparatively muted, so the patient may be floridly deluded and yet speak in a connected, followable way. Second, **catatonic phenomena** are rare in the epileptic psychosis, whereas catatonic features can arise in the course of primary schizophrenia.

These are, once again, tendencies rather than absolutes, and they have to be read alongside everything else rather than used as isolated rules. Their value is corroborative. A psychosis dominated by well-preserved, communicable positive symptoms, without significant formal thought disorder and without catatonia, in a patient who is affectively warm and cooperative, fits the epileptic picture. But because they are relative they can only tilt the balance; a patient with the epileptic psychosis may on occasion show some thought disorder, and their absence early in a primary schizophrenic illness proves nothing. The lesson that runs through every one of these discriminators is that each is probabilistic, and the diagnosis is a gestalt assembled from all of them.

Preserved	Non-progressive	Absent	Rare
AFFECT AND PERSONALITY	LONG-TERM COURSE	FORMAL THOUGHT DISORDER	CATATONIA

19.7 Putting the discriminators together

The whole is more than any part. No single feature listed above is pathognomonic, and any one of them can fail in an individual patient. The discrimination is made by assembling the features into a coherent picture and weighing that picture against the alternative. The comparison below sets the discriminating axes side by side; read down the epileptic column and a consistent gestalt emerges, a productive psychosis in a person whose affect, personality and course are comparatively intact, arising against a long history of seizures and without a schizophrenic family background.

DISCRIMINATING AXIS	INTERICTAL PSYCHOSIS OF EPILEPSY	PRIMARY SCHIZOPHRENIA
Affect	Relatively preserved, warm and responsive	Flattening early and marked
Interpersonal manner	Friendlier, cooperative, approachable	Guarded, suspicious, withdrawn
Premorbid personality	Relatively well preserved	Often abnormal or declining before onset
Long-term course	Non-deteriorating	Tendency to progressive deterioration
Family history	No excess familial schizophrenia	Increased familial incidence
Relation to seizures	Follows years of established epilepsy	No such antecedent
Formal thought disorder	Often absent	May be prominent
Catatonia	Rare	May occur

The discriminating axes read as a single gestalt; no row is pathognomonic on its own, and the diagnosis is made by the consistency of the column rather than by any one cell.

MNEMONIC

SPARE - the person is spared

The features that keep the diagnosis on the epileptic side spell **SPARE**. **S**eizures precede the psychosis by years. **P**remorbid personality and function preserved. **A**ffect warm and responsive. **R**elationships cooperative rather than suspicious. And no progressive d**E**terioration of course. To these the phenomenology adds positive symptoms without formal thought disorder or catatonia, and the pedigree adds no increased familial incidence of schizophrenia.

RULE **The psychosis of epilepsy is distinguished from schizophrenia longitudinally and contextually, not cross-sectionally.** Because the cardinal symptoms overlap, a single mental state examination cannot settle the question; the discriminating information lives in the history and in the course. The sequence of events and the trajectory over time outweigh any cross-sectional symptom.

19.8 Why the distinction earns its keep

The label changes prognosis, counselling and treatment. Making this differentiation is not an academic exercise, because the diagnosis alters three things at once. Prognostically, the preserved personality and the non-deteriorating course mean the outlook is generally better than in primary schizophrenia, and the patient should not be burdened with the expectations that attach to the primary illness. For counselling, the absence of a schizophrenic familial loading changes what relatives are told about their own risk. For treatment, recognising that the psychosis is a complication of epilepsy directs attention both to optimising seizure control and to choosing an antipsychotic with the seizure threshold in mind. The specific question of antipsychotic selection in epilepsy, and the mechanism by which antipsychotics lower the seizure threshold, are handled in their own right elsewhere in this chapter and are only flagged here.

In practice, a small number of questions settle most cases:

- Is there a well-established epilepsy of long standing that preceded the psychosis?
- Are the affect, personality and interpersonal manner relatively preserved?
- Is the long-term course free of the progressive deterioration of the primary illness?
- Is the family history free of a schizophrenic loading?

A final caution frames the whole topic. Every discriminator here is relative and probabilistic, and none is a rule that holds in every case; the diagnosis is the weight of evidence, not a single decisive sign. The differentiation is also specific to the chronic interictal presentation. The ictal psychosis that is a manifestation of the seizure itself, and the postictal psychosis that follows a cluster of seizures after a lucid interval, are separated from primary schizophrenia mainly by their timing and are dealt with as distinct entities. Hold the interictal discriminators for the illness they were built to describe: the chronic, schizophrenia-like psychosis that arises between seizures in

long-standing epilepsy and imitates the primary disorder closely enough to make the question worth asking in the first place.

DO NOT MISS

One organic mimic must never be missed at the other end of the timeline. A new psychosis with seizures, especially one that is acute or subacute in a young person, can be **anti-NMDA receptor encephalitis**: it classically presents in a young woman with a viral-like prodrome, a psychiatric-first picture, seizures, orofacial and limb dyskinesias, autonomic instability and a falling level of consciousness; it carries an association with an ovarian teratoma; and it is confirmed by paired serum and cerebrospinal fluid testing, with cerebrospinal fluid anti-GluN1 (NR1) IgG the decisive test because serum alone gives both false positives and false negatives. Treatment is immunotherapy together with a tumour search and resection of a teratoma when one is found, since many patients have no tumour. Its acute, dyskinetic, autonomically unstable course is the temporal opposite of the chronic interictal psychosis this topic discriminates, which is exactly why a fresh psychosis with seizures should trigger the autoimmune-antibody workup before it is filed as an epilepsy psychosis.

IN INDIA

The discrimination leans heavily on history and phenomenology, and that is exactly what makes it valuable in Indian practice. Where the epilepsy treatment gap is wide and many patients reach services only after years of poorly controlled seizures, the schizophrenia-like psychosis is a real and under-recognised presentation, and it is easily mislabelled as primary schizophrenia when longitudinal records, prior electroencephalography or even a reliable seizure history are unavailable. The clinician who cannot fall back on investigations must lean on the bedside discriminators, the preserved affect and personality, the non-deteriorating course, and above all the sequence of epilepsy first and psychosis later, and must take an active seizure history rather than waiting for one to be offered.

Preserved affect and personality, a course that does not deteriorate, a family history free of schizophrenia, and a psychosis that appears only after years of established epilepsy: these place a schizophrenia-like presentation on the epileptic side, even when the cross-sectional symptoms are indistinguishable.

20 · Antipsychotic choice and seizure-threshold considerations in epileptic psychosis

The antipsychotic is chosen against the seizure, not against the psychosis. Antipsychotic choice in epilepsy is a seizure-threshold consideration before it is a matter of psychiatric efficacy. In a patient whose psychotic episode arises on a background of epilepsy, the drugs that treat the psychosis are the same drugs that can destabilise the very seizures the neurologist has worked to control, and the whole skill of the topic lies in resolving that tension. The positive symptoms of an epileptic psychosis, whether it is a postictal psychosis, the schizophrenia-like psychosis of epilepsy or a psychosis emerging with forced normalisation, respond to antipsychotics much as psychosis does anywhere else. What differs, and what the marks reward, is the second question

the epilepsy forces on you: which of the effective agents is least likely to lower the seizure threshold and tip a vulnerable brain back into fits. Candidates who answer only the first question, prescribing as though the patient had no epilepsy, miss the point of the question entirely.

Two boundaries keep this topic clean. The property that ranks the agents, the dose-dependent lowering of the seizure threshold and clozapine's place at the extreme of it, is the substance of its own dedicated topic; here that ranking is taken as given and applied to a choice. And clozapine's full haematological and cardiac profile, with its monitoring, belongs to the Schizophrenia: management, antipsychotics and adverse effects notes. What is owned here is the decision itself: which agent, at what dose, introduced how, and with what antiepileptic cover.

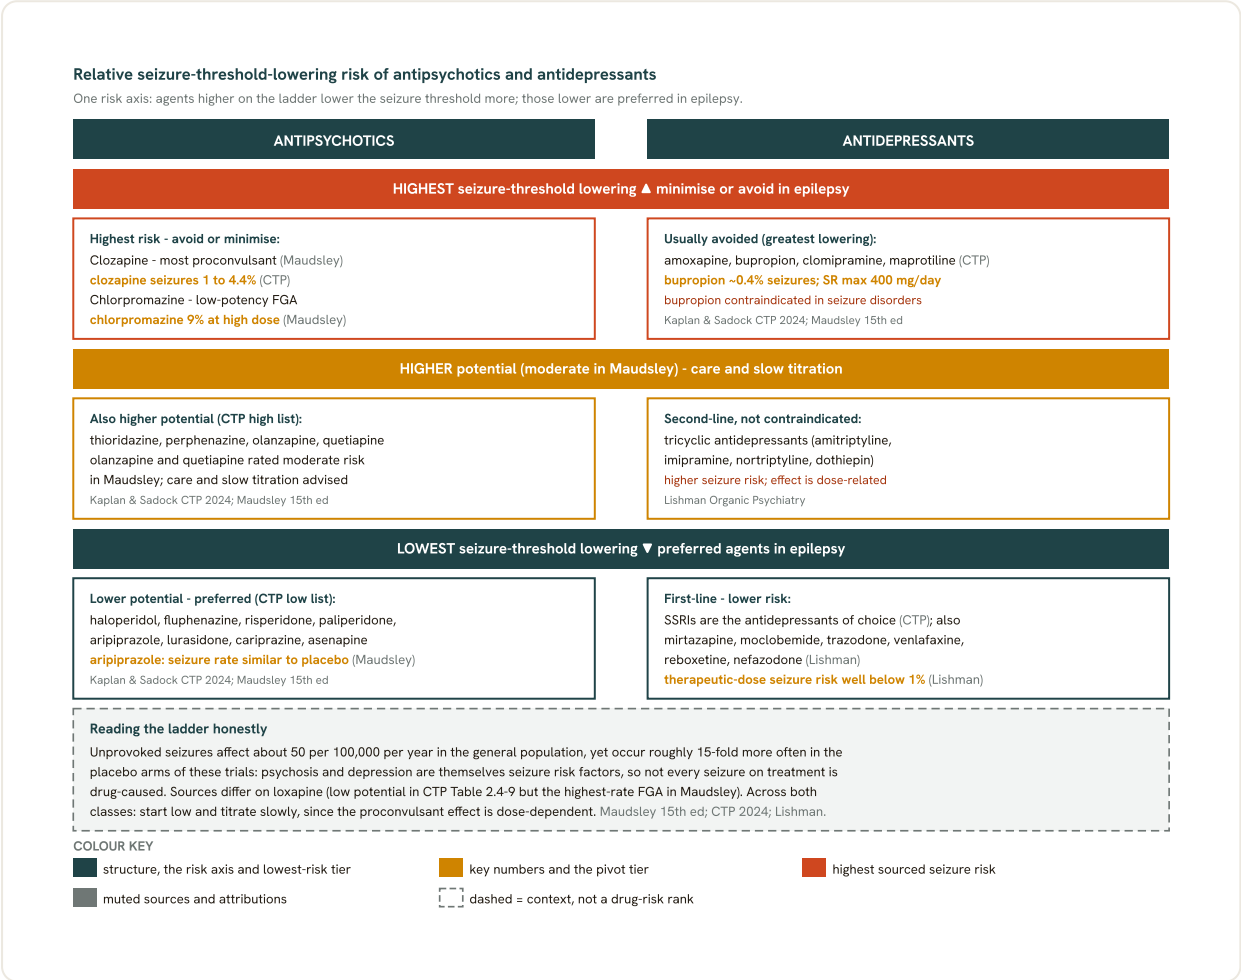


Fig 20.1 Relative seizure-threshold-lowering risk of antipsychotics and antidepressants. On a single risk axis, clozapine (the most proconvulsant antipsychotic, seizures in 1 to 4.4 per cent), high-dose chlorpromazine (9 per cent above 1 g per day) and the usually-avoided antidepressants amoxapine, bupropion, clomipramine and maprotiline sit highest, whereas high-potency first-generation antipsychotics, risperidone, paliperidone, aripiprazole and the SSRIs lower the threshold least and are preferred in epilepsy. Attributions per box (Kaplan and Sadock CTP 2024, Maudsley 15th ed, Lishman).

20.1 Why the choice is a threshold decision, not an efficacy contest

Efficacy is a wash; safety on the seizure threshold is the whole question. The starting fact, and the one candidates most often get wrong, is that there is no convincing difference in antipsychotic efficacy between the **first-generation antipsychotics** and the **second-generation antipsychotics** when they are used to treat the psychoses of epilepsy. If both classes bring the psychosis under control to a broadly similar degree, then efficacy cannot be the variable that

decides the prescription. The variable that can is the agent's propensity to lower the **seizure threshold**, and it is precisely this that becomes the primary selection criterion. The reasoning is worth stating explicitly because it inverts the habit of general adult psychiatry, where the choice usually turns on the efficacy and side-effect profile for the psychosis alone. In the epileptic patient the question is reframed: not which antipsychotic is most effective, but which effective antipsychotic is least likely to provoke a seizure in a brain that is already predisposed to them. Everything else in the topic, the preferred agents, the dosing, the method of introduction and the special problem of clozapine, is downstream of this single reorientation. The mechanism that generates the ranking, why some agents disturb the threshold more than others and why the effect is dose-related, is developed in its own topic and is not the subject here; the working clinician needs the ranking itself and the prescribing discipline that follows from it.

RULE **Select the antipsychotic by its effect on the seizure threshold, then introduce it low and slow.** Because the agents do not differ meaningfully in how well they treat the psychosis, the choice is made on seizure safety, and even the safest agent is made safer still by starting at a low dose and titrating gradually rather than loading.

20.2 The preferred agents at the safe end of the ranking

A small set of agents disturbs the threshold least and forms the default choice. The antipsychotics that lower the seizure threshold least, and are therefore preferred as first-line agents in a patient with epilepsy, are haloperidol, risperidone, paliperidone and aripiprazole. This is the set to reach for when an epileptic psychosis needs treating and there is no overriding reason to do otherwise. Among them, **aripiprazole** deserves separate emphasis, because it not only rarely lowers the seizure threshold but produced a seizure incidence comparable to placebo in randomised controlled trials, which is about as reassuring a signal as the evidence offers for any antipsychotic in this setting. That safety record is why it recurs on every low-risk list. The practical message is that a clinician does not need an exotic agent to prescribe safely here: two widely available high-potency options and two commonly stocked second-generation agents cover most situations, and their standing as preferred choices rests on seizure safety rather than on any claim to superior antipsychotic effect. Holding this short list in mind converts an anxious, open-ended decision into a quick one, which matters when the psychosis is florid and a drug must be started promptly.

MNEMONIC

HARP

The four preferred, least threshold-lowering antipsychotics in epilepsy. **H**aloperidol, **A**ripiprazole, **R**isperidone, **P**aliperidone. Each is an agent this topic names as a first choice on seizure-safety grounds; reach inside this set before you look outside it.

20.3 Grading the field: low risk against use with care

A working ranking sorts the antipsychotics into a safe tier and a care-required tier. The four preferred agents sit inside a slightly broader low-risk grouping that a clinician can draw from, and

it is useful to see that group set against the agents that call for caution. The Maudsley Prescribing Guidelines provide exactly this contrast in their antipsychotic table for people with epilepsy. On the safe side sit the low-risk good choices; on the other, agents at **moderate seizure risk** where more care is required, with a further pharmacokinetic wrinkle for quetiapine. The accompanying table lays the two tiers side by side.

RISK TIER	ANTIPSYCHOTICS	PRACTICAL NOTE
Low risk (good choices)	amisulpride or sulpiride, aripiprazole, ziprasidone, high-potency first-generation agents such as haloperidol and fluphenazine, and risperidone	preferred where an antipsychotic is genuinely needed
Care required (moderate risk)	olanzapine and quetiapine	quetiapine additionally carries a high risk of interaction with antiepileptic drugs

A practical two-tier grouping of antipsychotics for people with epilepsy: the low-risk agents to prefer, and the two that can be used but demand closer attention to dose and seizure control.

Reading the table the right way matters more than memorising it. The low-risk tier is not a licence to ignore the threshold; it identifies the agents least likely to cause trouble when the prescribing discipline described below is followed. The care-required tier does not forbid its agents either; olanzapine and quetiapine can be used, but they demand closer attention to dose and to seizure control, and quetiapine's high liability for interaction with antiepileptic drugs layers a pharmacokinetic hazard on top of the pharmacodynamic one. The agents that sit at the genuinely seizure-provoking end of the ranking, with clozapine the paradigm, are handled in their own topic and in the special case set out below, because they change the whole calculus of prescribing rather than sitting on a simple scale of caution.

20.4 Introducing the drug: low, slow and to a defined dose target

How the drug is started matters as much as which drug it is. Even a low-risk agent can precipitate a seizure if it is introduced carelessly, because rapid introduction and high dosing are the principal provokers of fits with these medicines. The prescribing rule is therefore to begin at a low dose and titrate upward slowly, letting the patient settle at each step rather than pushing to a target in a hurry. The dose that is actually needed is conventionally expressed in **chlorpromazine-equivalent** terms, which lets agents of very different milligram potencies be compared on a common scale. For a mild-to-moderate psychotic episode the usual requirement is **300 to 1,000** mg/day chlorpromazine-equivalent, reached by that same gradual titration rather than by front-loading, and in epilepsy the aim is the lower end of that range that still controls the psychosis. Dosing above the usual range is exceptional high-dose treatment, resorted to only when standard doses have genuinely failed and only under specialist supervision with intensified antiepileptic-level and EEG monitoring, because it is at the high end that seizure liability climbs. Severity of illness does not by itself license a higher target: the figures describe a range, not a licence to push the dose up quickly. The same total dose is far safer approached in small increments than imposed at once, which is the entire reason the method is emphasised in the

same breath as the number: a correct target reached by an incautious route can still provoke the seizure the choice of agent was meant to avoid. The dose range is also a reminder that undertreatment is its own failure; a psychosis left under-dosed out of excessive threshold anxiety is not managed, merely neglected, and the answer is a safe agent taken to an adequate dose by a safe route rather than a good agent kept ineffectually low.

20.5 Prescribing under surveillance: antiepileptic levels and the EEG

The titration is carried out watched, not blind. Adding an antipsychotic to an established antiepileptic regimen is a two-way pharmacological event, and it is monitored on two fronts. As the antipsychotic is introduced and titrated, antiepileptic drug levels and the EEG should both be followed. The drug-level check exists because antipsychotics and antiepileptics interact: enzyme induction or inhibition can move an antiepileptic concentration out of its effective range and undo seizure control silently, and the detailed pharmacokinetics of those enzyme-induction and inhibition interactions are set out in their own topic rather than reopened here. The EEG check exists because the threshold effect can announce itself electrographically before it declares itself as a convulsion, so a worsening record is an early warning that the balance is tipping while there is still time to act. Together the two turn an otherwise blind escalation into a monitored one, and they give the clinician somewhere concrete to look when deciding whether a rising dose is safe to continue or whether it has begun to cost seizure control. The same surveillance frames the response to a breakthrough seizure: it is not automatically a reason to abandon the antipsychotic, but it is a reason to reconsider the whole balance of the regimen.

- Pick an agent from the low-risk tier unless there is a compelling reason not to.
- Start below the eventual target and titrate upward slowly.
- Follow antiepileptic drug levels as the antipsychotic is added.
- Watch the EEG for early threshold change before it becomes a fit.

20.6 When clozapine is unavoidable: the antiepileptic pairing

The one agent you least want in epilepsy is sometimes the only one that works. Treatment-resistant psychosis can leave **clozapine** as the only effective option even in a patient with epilepsy, and although clozapine sits at the seizure-provoking end of the ranking, refusing it outright may not be tenable when nothing else has worked. Its position on that ranking and the mechanism behind it are the business of the dedicated seizure-threshold topic, and its full haematological and cardiac profile with the associated monitoring belongs to the Schizophrenia: management, antipsychotics and adverse effects notes; neither is re-taught here. What this topic owns is the co-prescribing decision, and it has a clear answer. When clozapine must be used in epilepsy, **lamotrigine** is the antiseizure medication of choice to pair with it, and **carbamazepine** should be avoided. Carbamazepine is the wrong partner for two converging reasons: in combination it raises the risk of blood dyscrasias, stacking a haematological hazard onto an agent that already demands haematological vigilance, and it lowers clozapine levels, threatening to undercut the very efficacy that forced the choice of clozapine in the first place. Lamotrigine's own antiepileptic titration and rash are taught in their own topic, and its separate mood-stabiliser role

belongs to the Mood disorders: pharmacotherapy notes; here it is simply the safer antiepileptic companion for a difficult but sometimes unavoidable antipsychotic.

PITFALL

A common error when starting clozapine in a patient with epilepsy is to reach for carbamazepine as the antiepileptic cover, on the loose reasoning that it is also a familiar mood stabiliser and will do double duty. That choice doubles the blood-dyscrasia hazard and simultaneously drops clozapine levels, degrading safety and efficacy at the same time. Pair clozapine with lamotrigine instead, and keep carbamazepine well away from it.

20.7 Putting the choice into a management sequence

What to actually do, and in what order. The prescribing decision sits inside a broader management logic that is anchored, for the epilepsy itself, in the ILAE and NICE epilepsy guidelines, and for the psychiatric prescribing, in the epilepsy-specific psychiatric-management literature and the Maudsley Prescribing Guidelines. The governing principle comes first and is easy to state: optimise seizure control before any psychotropic is started, and if a seizure occurs after one has been introduced, either change the psychotropic or optimise the antiseizure medication. Before committing to an antipsychotic at all, the reversible and mimicking causes of the disturbance should be excluded, because a peri-ictal or iatrogenic psychosis is treated by addressing its cause rather than by adding a threshold-lowering drug on top of it.

The **LOCUS** of care follows the severity and the risk. A stable patient with a mild presentation can be managed as an outpatient with careful titration and review; a florid psychosis carrying behavioural risk warrants inpatient admission where dosing and seizures can be watched closely; an acutely disturbed patient may first present to an emergency setting. The **focus** shifts by phase. The acute target is control of the psychosis with the least threatening effective agent at a low starting dose. The short-to-mid-term target is titration to the chlorpromazine-equivalent dose the severity requires, carried out under antiepileptic-level and EEG surveillance. The long-term target is durable seizure control, which is the best protection against both further fits and recurrence of the psychosis, together with periodic review of whether the antipsychotic is still needed at all. The **modus** spans every mode of treatment. The pharmacological mode is the core of this topic: an agent chosen from the low-risk tier, introduced low and slow, with clozapine reserved and specifically paired with lamotrigine when it cannot be avoided. The procedural mode is the monitoring itself, the antiepileptic levels and the EEG that make a safe escalation possible, alongside optimisation of the antiepileptic regimen as the durable lever on the disorder. The social and psychological modes carry the rest: securing adherence, educating the patient and family about the interaction between the two treatments, and coordinating with the treating neurologist so that seizure control and psychiatric control are not pursued at cross purposes.

IN INDIA

Access shapes the choice as much as pharmacology does. Several of the preferred and low-risk agents, haloperidol and risperidone in particular, are inexpensive, widely stocked and familiar across public and private settings, which makes safe prescribing achievable even where budgets are tight; newer agents such as aripiprazole and paliperidone are available but cost more, so the low-cost members of the low-risk tier often become the practical default. Clozapine can be used but demands the haematological monitoring infrastructure that is easier to guarantee in tertiary centres than in district practice, so the threshold for choosing it, already high on seizure grounds, is higher still where reliable monitoring is hard to sustain. The pragmatic answer, a low-cost high-potency or low-risk second-generation agent, started low and titrated slowly with close attention to seizure control, fits Indian practice well and rarely requires anything exotic.

GOOD TO REMEMBER

When an epileptic psychosis needs an antipsychotic, default to one of the preferred agents, start at a low dose and titrate slowly, keep the seizures optimally controlled first, and monitor antiepileptic levels and the EEG as you go. Do not assume a second-generation drug is automatically the safer choice on the threshold, and never pair clozapine with carbamazepine.

TRAP **Assuming newer means safer on the seizure threshold.** Candidates reflexively pick a second-generation antipsychotic as the seizure-safe option, but efficacy does not separate the classes and some second-generation agents are worse on the threshold, while a high-potency first-generation agent such as haloperidol is among the safest. The generation of the drug does not predict its threshold liability; its specific place in the ranking does.

300 to 1,000	EEG + levels	near placebo
MG/DAY CPZ-EQUIVALENT, MILD TO MODERATE EPISODE	MONITOR DURING ANY ESCALATION, ABOVE ALL AT HIGH DOSE	ARIPIPRAZOLE SEIZURE INCIDENCE IN TRIALS

In epileptic psychosis the antipsychotics do not differ meaningfully in efficacy, so choose on the seizure threshold: prefer haloperidol, risperidone, paliperidone or aripiprazole, start low and titrate slowly under antiepileptic-level and EEG monitoring, optimise seizure control first, and if clozapine is unavoidable pair it with lamotrigine and avoid carbamazepine.

21 · Antipsychotic-induced lowering of seizure threshold (clozapine and others)

The drug that quiets the psychosis can uncover the seizure. Every clinician who starts an antipsychotic in a person with epilepsy, or in a person whose psychosis has arisen out of epilepsy, works against a built-in tension: the molecules that suppress positive symptoms also nudge the cortex toward its convulsive threshold. This topic is the mechanism behind that tension. The phenomenon is antipsychotic-induced lowering of the seizure threshold, and its hallmark is a dose-dependent seizure liability. It sets out what it means to say a drug lowers the **seizure threshold**, how far the effect depends on how the drug is given rather than on the drug alone, and which agents carry the liability most heavily. It earns marks because it is the single pharmacological principle that governs a whole cluster of downstream prescribing decisions, and

because the figures attached to the worst offenders are exactly the kind of discrete fact an examiner can ask for by name. The seizure threshold is a shorthand for the level of excitability at which a synchronised discharge becomes self-sustaining; anything that tips the balance of cortical excitation and inhibition can lower it, and most antipsychotics do so to some degree.

Two boundaries frame the discussion before it opens. First, this is the mechanism, not the prescription: which antipsychotic to actually select for a patient with epileptic psychosis, and how the ranking is weighed at the bedside, is developed in its own topic and is not reopened here. Second, where clozapine is concerned this account covers only its effect on the seizure threshold; its agranulocytosis, myocarditis and the blood monitoring that surrounds it belong to the Schizophrenia: management, antipsychotics and adverse effects notes. Keeping those two lines drawn stops the topic sprawling into the two that flank it and keeps the mechanism itself in sharp focus.

21.1 What lowering the seizure threshold actually means

A graded, dose-driven effect, not an on-off switch. The central fact of the whole topic is that antipsychotics lower the seizure threshold in a **dose-dependent** way, with the danger greatest when the drug is introduced quickly and driven to high doses. The distinction that carries the marks is that proconvulsant liability is not a fixed, all-or-nothing property of a molecule. The same agent given at a low dose reached slowly may sit almost inertly at the threshold, while the identical agent pushed rapidly to a high dose becomes meaningfully epileptogenic. Mechanistically the effect reflects a shift in the working balance between cortical excitation and inhibition, and it shows itself on the electroencephalogram before it shows itself as a fit: many antipsychotics produce dose-related background slowing, and at higher concentrations some provoke frank epileptiform discharges in susceptible brains. Kaplan and Sadock's Comprehensive Textbook of Psychiatry frames the whole class this way, as a graded liability rather than a categorical one, and the practical corollary is the reason the topic matters. If the effect scales with dose and with the speed of introduction, then a meaningful share of the risk is created by the way the drug is prescribed and not only by the drug that was chosen. That is what turns an apparently fixed pharmacological hazard into something the clinician can partly control. It also explains why two prescribers using the same agent in the same patient can generate quite different seizure risk, and why a blanket statement that a given antipsychotic is safe or unsafe, stripped of any reference to dose and titration, is close to meaningless.

21.2 The determinants you can change, and the ones you cannot

Rate and concentration are the modifiable levers. Drilling into the dose-dependence gives the practically useful version of the principle. The two determinants that a clinician can move are **rapid dose escalation** and the resulting high plasma concentration, and these are where the modifiable risk lives. Everything else is comparatively fixed: the intrinsic proconvulsant ranking of the particular agent, and the patient's own underlying epilepsy or predisposition. Reading the risk this way, split into what can be changed and what cannot, is what lets the pharmacology be turned into safer prescribing without needing any further numbers.

- Modifiable: how fast the dose is escalated, since a slow, stepwise build carries less risk than an abrupt loading schedule.
- Modifiable: the target dose and the plasma level it produces, since risk climbs with concentration.
- Fixed: the intrinsic liability of the chosen agent, which sets the floor of risk before any titration decision.
- Fixed: the patient's own seizure disorder, which is the substrate the drug acts upon.

RULE **Rate and dose, not the drug alone.** The seizure-threshold effect of an antipsychotic is dose-dependent and is worst with rapid introduction at high dose. The single safest manoeuvre with any agent is therefore to titrate slowly to the lowest effective dose, which is a lever that applies before the question of which molecule to use is even asked.

MNEMONIC

Dose, Rate, Drug, Disease

Four things set the seizure risk of an antipsychotic, and each is a fact this topic states. **D**ose - the higher the dose, the greater the risk. **R**ate - the faster the escalation, the greater the risk. **D**rug - clozapine leads the field for proconvulsant liability. **D**isease - a poorly controlled epilepsy is where the effect stops being academic.

21.3 The attribution problem: not every seizure is the drug

A seizure on treatment is not proof the treatment caused it. Before any antipsychotic is blamed for a fit, a confound has to be reckoned with, and it is one of the more sophisticated points the topic can be asked to make. The background incidence of **unprovoked seizures** in the general population runs at about **50 per 100,000 person-years**, but in the placebo arms of antipsychotic and antidepressant trials the seizure rate is roughly **15-fold higher** than that baseline, as the Maudsley Prescribing Guidelines set out. The people being treated with these drugs are, in other words, already at elevated risk before a single tablet is swallowed, because psychosis and depression are themselves seizure risk factors rather than neutral backgrounds. The implication for attribution is direct and clinically important: a proportion of the seizures that occur in patients on antipsychotics would have occurred anyway, driven by the illness and its correlates rather than by the medication. This does not exonerate the drugs, whose dose-dependent contribution is real, but it does mean that a seizure appearing during treatment cannot be read as automatic proof of drug causation. The honest position is that the drug adds to a risk that was already raised, and separating the two contributions in an individual patient is genuinely difficult.

PITFALL

The reflex when a patient has a seizure on an antipsychotic is to indict the drug and stop it. Because psychosis itself raises seizure risk well above the population baseline, that reflex over-attributes: some of these seizures would have happened without the medication, and abruptly withdrawing an effective antipsychotic to chase a seizure it may not have caused can leave the patient both psychotic and no safer. Weigh the dose, the titration speed and the underlying illness before pinning the event on the drug.

21.4 The proconvulsant ladder

A ranking from high liability to low. The antipsychotics can be arranged along a gradient of seizure-threshold-lowering potential, and the accompanying figure lays this ladder out from the highest-risk agents down to the lowest. Kaplan and Sadock's table gives the two ends explicitly. The high-potential tier gathers chlorpromazine, clozapine, thioridazine, perphenazine, olanzapine and quetiapine, agents that between them include the low-potency phenothiazines and the dibenzo-structured compounds. The low-potential end gathers fluphenazine, haloperidol, risperidone, paliperidone, aripiprazole, lurasidone, cariprazine and asenapine, along with loxapine, molindone and pimozide, a group weighted toward the high-potency first-generation drugs and most of the newer second-generation agents.

TIER ON THE LADDER	ANTIPSYCHOTICS	STRUCTURAL FLAVOUR
High seizure-threshold-lowering potential	chlorpromazine, clozapine, thioridazine, perphenazine, olanzapine, quetiapine	low-potency phenothiazines and dibenzo compounds cluster here
Low seizure-threshold-lowering potential	fluphenazine, haloperidol, risperidone, paliperidone, aripiprazole, lurasidone, cariprazine, asenapine (with loxapine, molindone and pimozide)	high-potency first-generation agents and most newer drugs cluster here

The seizure-threshold-lowering ladder as set out in Kaplan and Sadock's table, from highest to lowest potential; the tiers, not a precise rank order within each tier, are what carry the teaching point.

The ladder is worth committing to memory as a shape rather than as a rote list, because its two ends behave differently and because the placement of certain agents is exactly where the sources part company. The structural pattern helps: the drugs at the high end are dominated by the sedating, low-potency and tricyclic-shaped compounds, while the high-potency agents and most of the newer molecules sit lower. That heuristic is a scaffold for recall, not a substitute for the table, and the exceptions to it are the substance of the next two sections.

21.5 Clozapine: the agent at the top

The most proconvulsant antipsychotic, and dose tells the tale. If a single agent has to be named as the seizure-threshold offender, it is clozapine, which the Maudsley Prescribing Guidelines identify as the **most proconvulsant** antipsychotic in use. The quantitative anchor is that clozapine induces seizures in **1% to 4.4%** of the patients who take it, and the spread of that range is itself instructive: the risk sits toward the lower bound with careful, gradual dosing and

rises toward the upper bound when the dose is increased rapidly. Clozapine thus reproduces the whole topic in one molecule. It is the drug whose intrinsic liability is highest, and it is also the drug in which the dose-dependence and the danger of quick escalation are most vividly demonstrated, which is why its titration is conventionally slow and stepwise. Two cautions attach to teaching it. The first is a boundary: the seizure-threshold angle is all that is owned here, and clozapine's fuller adverse-effect profile, its haematological monitoring and its other hazards, belongs to the Schizophrenia: management, antipsychotics and adverse effects notes and should be cross-referred rather than repeated. The second is that a seizure on clozapine is not by itself a reason to abandon the drug in a patient who needs it; it is a reason to review the dose, the plasma level and the speed of the recent titration, since those are the modifiable determinants the mechanism has already identified.

21.6 Chlorpromazine and loxapine: the first-generation flags

Two older agents that punch above their apparent place. Among the first-generation antipsychotics two names deserve to be flagged in their own right. Chlorpromazine, the archetypal low-potency phenothiazine, carries a clear dose-threshold effect: doses above **1 g per day** are associated with a seizure incidence of **9%**, a figure high enough that high-dose chlorpromazine is best avoided in anyone with a seizure tendency. Loxapine is the more interesting teaching case, because it is where two authorities pull apart. On Kaplan and Sadock's ladder loxapine sits in the low-potential tier, yet the Maudsley singles it out as the first-generation antipsychotic carrying the **highest seizure rate** of its class. Both statements come from reputable sources and both are worth knowing; the resolution is not to decide which is right but to recognise that loxapine, structurally a dibenzoxazepine and so a chemical cousin of the high-risk dibenzo compounds, is an agent about which the sources genuinely disagree, and to treat it with corresponding caution. The general lesson the two drugs teach together is that structural family and dose, not the drug's marketing generation, predict seizure liability.

IN INDIA

Cost and availability keep the older low-potency antipsychotics, chlorpromazine foremost among them, in far heavier routine use in much of Indian practice than in settings where newer agents are the default. That matters directly for this topic, because it is precisely these low-potency phenothiazines that occupy the high end of the seizure-threshold ladder, and because high total daily doses are still encountered where a cheaper drug is pushed rather than a more expensive one substituted. The **9%** seizure incidence attached to chlorpromazine above **1 g per day** is therefore not a historical curiosity here but a live prescribing consideration. A related caution is that co-prescribed antiepileptic drugs alter antipsychotic plasma levels through enzyme effects, a pharmacokinetic interaction that is covered in its own topic and that further complicates the dose-driven risk in exactly the patients most likely to be on both.

21.7 The second-generation gradient and why the ladder is source-dependent

The ranking is a guide, not a law. The newer antipsychotics do not lower the seizure threshold uniformly, and the same discrepancy that surfaced with loxapine reappears among them, which is what makes this the most examinable subtlety of the topic. Olanzapine and quetiapine sit at the

high end of Kaplan and Sadock's ladder, yet the Maudsley, weighing the randomised-trial data, judges their overall proconvulsant risk to be **low-to-moderate**, with the olanzapine evidence frankly mixed, since both anticonvulsant and proconvulsant reports exist for it. The point is not that one source is mistaken. It is that seizure-threshold rankings are assembled from heterogeneous evidence, case reports, trial adverse-event data and pharmacological reasoning, and different authorities weight those sources differently, so the same agent can legitimately appear in a high tier on one ladder and a lower band on another. A candidate who has memorised a single ranking and recites it as though it were fixed will be caught out by exactly these agents. The defensible position is to know the shape of the ladder, to know its firm anchors, clozapine at the top and the high-potency first-generation drugs low down, and to treat the placement of olanzapine and quetiapine as genuinely source-dependent rather than settled.

TRAP **Reciting one ladder as gospel.** The classic error is to learn a single seizure-threshold ranking and answer as though every agent's place were fixed. The sources disagree at specific rungs: loxapine is low-potential on one ladder but the highest-risk first-generation drug on another, and olanzapine and quetiapine sit high on one and low-to-moderate on the other. Name the disagreement rather than picking a side, and anchor the answer on the points that do not move.

21.8 How much it matters in the patient with epilepsy

Usually subclinical, decisive in the poorly controlled. The final calibration is the one that stops the topic being alarmist. For most patients the seizure-threshold-lowering effect of an antipsychotic is real but not clinically significant; it is a measurable pharmacological property that never translates into an actual fit. Where it changes from a footnote into a genuine hazard is in the patient whose epilepsy is already **poorly controlled**, in whom the drug's contribution is added to a threshold that is already low and a brain already prone to discharge. Kaplan and Sadock make exactly this qualification, that the effect is usually subclinical but can reach clinical significance in the setting of poorly controlled epilepsy. That single distinction organises the whole clinical stance. In a person with no seizure history and well-controlled or absent epilepsy, the threshold effect rarely dictates practice on its own; in a person with active, poorly controlled seizures it becomes one of the real constraints on treatment, and it is the reason the choice of agent and its titration are handled with such care in that population. The mechanism laid out across this topic, dose-dependence, the modifiable levers of rate and concentration, the ladder with clozapine at its head, and the attribution confound, all converges on this practical endpoint: the same drug is a minor consideration in one patient and a major one in another, and it is the state of the underlying epilepsy that decides which.

1 to 4.4%	9%	~50/100k/yr	~15-fold
CLOZAPINE SEIZURE INCIDENCE	CHLORPROMAZINE ABOVE 1 G/DAY	BASELINE UNPROVOKED-SEIZURE RATE	HIGHER IN TRIAL PLACEBO ARMS

GOOD TO REMEMBER

When you add an antipsychotic to someone with epilepsy or a prior seizure, the most useful thing you control is not always the molecule but the manner: build the dose slowly to the lowest effective level, since the effect is dose-dependent and worst with rapid escalation. Keep clozapine and high-dose low-potency phenothiazines in mind as the high-liability end, remember that a seizure on treatment may reflect the illness as much as the drug, and leave the actual agent selection for the epileptic-psychosis patient to the topic that owns it.

Antipsychotics lower the seizure threshold in a dose-dependent way that is worst with rapid escalation and high doses; the effect is greatest with clozapine and with high-dose low-potency agents, is usually subclinical, and becomes clinically important mainly in poorly controlled epilepsy.

Mood, anxiety and suicide in epilepsy

22 · Depression in epilepsy: prevalence, phenomenology and screening

Depression is the disorder that most often accompanies epilepsy, and the one most often overlooked. Of all the psychiatric problems a person with epilepsy may carry, depression is at once the most frequent, the most damaging to the very quality of life that seizure treatment is meant to protect, and the most easily written off as an understandable reaction to a difficult illness. Together with anxiety it is one of the two most frequently encountered interictal psychiatric disorders of epilepsy, the anxiety half of that pair being developed in its own topic in this chapter. The examinable material falls into three blocks that map cleanly onto what a competent clinician actually does: knowing how common depression really is and why the published figures scatter so widely, recognising that its presentation is frequently atypical and therefore does not announce itself, and screening for it deliberately rather than waiting for the patient to complain.

This topic earns marks out of proportion to its apparent simplicity because almost every part of it runs against a clinician's first instinct. The rates are higher than intuition suggests, the phenomenology is stranger than the textbook portrait of major depression, and the condition erodes daily functioning more than the fits themselves, yet it is the fits that dominate the consultation. The wider two-way traffic between epilepsy and psychiatric illness, in which each raises the risk of the other, is set out in its own dedicated topic in this chapter and is assumed here; this topic narrows to the depression that arises in the interictal state, drawing on Lishman's Organic Psychiatry and Kaplan and Sadock's Comprehensive Textbook of Psychiatry.

22.1 The disorder that owns the mood picture

Commonest, and the single heaviest drag on quality of life. The reason depression sits at the head of this band is not only that it is frequent but that it is disproportionately disabling. In patients whose seizures continue, depression is the single most important predictor of poor

quality of life, weighing more heavily on how the patient actually lives than seizure frequency, seizure severity or the burden of medication. That fact reorganises the whole clinical priority list. Epilepsy services are built to chase seizure control, and seizure freedom is the outcome the clinic measures itself by, yet a patient who has been rendered seizure-free can remain profoundly unwell if a coexisting depression goes unrecognised, while a patient with occasional seizures and good mood may function far better than the seizure count predicts. The tragedy compounded onto this is that the depression is routinely neither recognised nor treated, precisely because its atypical form keeps it below the clinician's threshold of suspicion. The mood disorder that does the most damage is therefore also the one most likely to be missed, and that pairing is the core of why the topic is worth teaching and worth examining.

RULE **Depression is the commonest interictal psychiatric disorder in epilepsy and the strongest single determinant of poor quality of life, yet it is routinely under-recognised because it seldom looks like textbook major depression.** Every clinical move that follows, from active screening to treating mood as seriously as seizures, descends from that one sentence. Seizure control is necessary but not sufficient for a good outcome.

22.2 How common it is, and why the figures scatter

The rate you quote depends entirely on how depression was defined and where the sample came from. No single prevalence figure for depression in epilepsy is defensible, and a candidate who offers one is usually offering a wrong one. The reported rates swing across an enormous range, and the swing is not noise but information: it tracks three things that must be held apart. The first is the level of the phenomenon being counted, since **depressive symptoms** on a self-report scale are far more common than a diagnosable **major depression** confirmed at structured interview. The second is the sampling frame, which climbs from unselected community cohorts, through neurology and specialist epilepsy clinics, to the highly selected presurgical populations of patients with medically refractory temporal lobe epilepsy, each tier concentrating the sickest and most treatment-resistant patients and so lifting the rate. The third is the diagnostic instrument and the window of time it interrogates. The accompanying table lays these out so the ascending gradient is visible rather than mysterious. Read that way, the figures are coherent: symptom-level counts in the community sit low and wide, clinic-based major depression sits in the middle, and lifetime rates in presurgical refractory disease sit highest, with major depressive episodes across pooled series reaching **up to 20 per cent**. The lesson to carry into the exam is comparative and mechanistic, not a single memorised percentage.

POPULATION AND SETTING	HOW DEPRESSION WAS DEFINED	REPORTED RATE
Community-based samples	depressive symptoms (not a diagnosed disorder)	4 to 37 per cent
Specialist epilepsy clinics	ICD or DSM major depression	17 to 21 per cent
Presurgical, medically refractory TLE	**lifetime prevalence** of major depression	24 to 35 per cent
Mixed epilepsy series (point and period)	DSM major depression	3 per cent over the past month, rising to 22 per cent over the previous year
Pooled across studies	any depression	7.5 to 34 per cent

Prevalence of depression in epilepsy sorted by the definition used and the population sampled; the spread is driven by measure, definition and sampling, not by genuine disagreement about the disease.

MNEMONIC

The three P's

Depression in epilepsy is **P**revalent (the commonest interictal psychiatric disorder), **P**leomorphic (atypical, frequently short of standard criteria), and **P**icked up only if you screen (it is rarely volunteered and easily missed). Each of the three is a fact this topic actually establishes, and together they are the spine of a complete answer on prevalence, phenomenology and detection.

22.3 It is unipolar depression, not bipolarity

The mood disorder of epilepsy is depressive; bipolar disorder is not over-represented. A neat and examinable negative sits alongside the prevalence data, and it is worth stating explicitly because intuition and older limbic theorising both pull the other way. Given the temporolimbic disturbance that underlies so much of the psychiatry of epilepsy, a candidate might expect an excess of bipolarity, but the evidence does not support it: bipolar disorder is not over-represented in epilepsy, and even in the enriched presurgical series where every psychiatric rate is inflated, its prevalence has been reported at **below 1 per cent**. The practical consequence is that a patient with epilepsy who presents with low mood should be worked up as a unipolar depression until there is specific evidence otherwise, and the reflexive reach for a mood-stabiliser framing on the strength of the epilepsy alone is unjustified. The corollary, that peri-ictal and interictal mood phenomena can nonetheless look labile and episodic, belongs to the phenomenology below and to the separate peri-ictal topics of this chapter, and must not be confused with a primary bipolar diathesis.

22.4 Depression tracks seizure control

In community samples, depression rises steeply as seizure control worsens. One of the most robust and most useful findings in this area is a dose-like relationship between how well the epilepsy is controlled and how much depression the population carries. In community-based work using the **Hospital Anxiety and Depression Scale**, on which each of the two seven-item subscales runs **0 to 21** with 0 to 7 read as non-case, 8 to 10 as borderline and 11 to 21 as a probable case, the proportion with significant depression climbs across bands of worsening

seizure frequency, and a second independent community series reproduced the same gradient with different absolute figures, which is exactly the kind of replication that turns an association into a fact worth teaching. The gradient makes intuitive sense: continuing seizures mean continuing unpredictability, injury, restriction and loss of autonomy, and the reactive burden mounts accordingly, while the same uncontrolled disease may also reflect the more extensive brain disturbance that independently generates mood symptoms. There is an important qualification, however, and it is examinable. The gradient is a feature of community and unselected samples; in tertiary referral populations it is not seen, because such samples are so concentrated with refractory, severely affected patients that the range of seizure control collapses and the relationship can no longer show itself. The accompanying table gives the community gradient.

SEIZURE CONTROL	SIGNIFICANT DEPRESSION
Seizure-free	4 per cent
Fewer than monthly seizures	10 per cent
More frequent seizures	21 per cent

Significant depression by seizure frequency in a community series (Hospital Anxiety and Depression Scale); a second community series reproduced the same rising gradient with figures of about 6, 11 and 33 per cent. The gradient is not seen in tertiary referral samples.

22.5 Who is most at risk

Focal temporolimbic epilepsy and poor seizure control carry the highest mood-disorder rates. Beyond the seizure-control gradient, a set of correlates identifies the patients in whom depression is most likely and in whom the clinician should look hardest. The strongest structural association is with **temporolimbic** epilepsy: focal impaired-awareness seizures of mesial-temporal origin, together with poor seizure control, are associated with higher rates of mood disorder. A left-hemisphere focus has been linked to greater depression, but this must be stated with the caution the source literature attaches to it, namely that any laterality effect is small and inconsistent, so it is a soft correlate rather than a rule to hang a diagnosis on. The seizure-type and localisation details themselves are taught in the temporal lobe and seizure-classification topics of this chapter and are named here only as risk markers. The correlates are best held as a short list, because in practice they braid together and each points to a different lever in formulation.

- Poor seizure control, the single most consistent and modifiable correlate.
- A temporolimbic, mesial-temporal focal onset, the strongest structural association.
- A left-hemisphere focus, a weak and inconsistent correlate not to be overstated.
- Antiepileptic drug treatment itself, an iatrogenic contribution whose detail belongs to the topic on drug-attributable mood and suicidality in this chapter.
- The psychosocial weight of stigma, unpredictability and lost independence, the reactive strand that no drug will reach.

22.6 Phenomenology: atypical and pleomorphic

It rarely reads as textbook major depression. The defining clinical feature of depression in epilepsy, and the one most likely to be tested, is that its presentation is frequently **atypical** and **pleomorphic**, so that a large proportion of cases fall short of the standard diagnostic criteria for a major depressive episode. Estimates that some **30 to 70 per cent** of depressive presentations in epilepsy show atypical features capture how far this departs from the classical picture. Instead of the sustained, criterion-complete syndrome of persistent low mood, anhedonia and neurovegetative disturbance, the mood disturbance of epilepsy is more often intermittent and fluctuating in its course, coloured by prominent irritability, and marked by endogenous somatic symptoms that shift and recur rather than settling into a single sustained episode. Because the syndrome is broken up and irregular, it slips through diagnostic instruments built to detect a continuous two-week episode, and it is easy for both patient and clinician to treat each fragment as a passing bad patch rather than a disorder. There is a recognised, named syndrome that formalises this atypical, chronic and pleomorphic interictal mood picture, interictal dysphoric disorder, but it is developed in full in its own dedicated topic in this chapter and is named here only to mark that the atypicality has been given a distinct diagnostic identity; its criteria and mechanism are not reopened here.

TRAP **Expecting the textbook major depressive episode and dismissing the diagnosis when it is not there.** Candidates and clinicians alike anchor on the DSM or ICD picture of a sustained two-week syndrome, and because the depression of epilepsy is intermittent, irritable and criterion-incomplete, they conclude the patient is not depressed. The examinable point is the opposite: an atypical, fluctuating, irritable mood disturbance that does not meet full criteria is the characteristic form here, not evidence against the diagnosis.

PITFALL

A common clinical error is to attribute the patient's low mood, fatigue, poor concentration and irritability to the epilepsy itself, to the sedating effect of antiepileptic drugs, or to the understandable distress of living with fits, and therefore to observe rather than treat. Each explanation may hold a grain of truth, but used as a reason not to act it leaves a treatable and disabling disorder to run. Name the depression before you explain it away.

22.7 Screening and detection

Because it is missed, it must be screened for, and there is an epilepsy-specific tool to do it. Everything above converges on a single practical conclusion: a disorder this common, this disabling and this atypical will not be caught by waiting for the patient to volunteer it, so detection has to be active and systematic. Two features make routine screening especially valuable in epilepsy. The atypical, intermittent phenomenology means the diagnosis is easily overlooked in an ordinary consultation dominated by seizures and medication, and the somatic overlap between depressive symptoms, seizure aftermath and drug side effects muddies clinical judgement. A structured screen cuts through both. The instrument developed for exactly this purpose is the **Neurological Disorders Depression Inventory for Epilepsy**, a brief screening tool designed for and validated in people with epilepsy, and constructed so that its items are not

confounded by the somatic and cognitive effects of seizures and antiepileptic drugs in the way generic depression scales can be. It is a six-item scale scored **6 to 24**, and a total **above 15** is a positive screen for major depression, a threshold chosen to be highly specific, at about 90 per cent, at acceptable sensitivity. A positive screen is a prompt to a full diagnostic assessment, not a diagnosis in itself. Once depression is confirmed, two onward questions arise that this topic hands to their owners: which antidepressant to choose against the drug's effect on the seizure threshold, which is the subject of its own dedicated topic in this chapter, and the assessment of the heightened suicide risk that depression in epilepsy carries, which likewise has its own dedicated topic. This topic ends at reliable detection, which is where the marks for depression in epilepsy are concentrated.

GOOD TO REMEMBER

Screen every patient with epilepsy for depression at routine review, not only those who look low, because depression predicts how the patient lives more powerfully than seizure frequency does. A seizure-free patient can still be doing badly, and the only way to know is to ask in a structured way rather than to infer mood from the seizure chart.

IN INDIA

In high-volume neurology and psychiatry clinics, where a consultation may last only a few minutes and is understandably consumed by seizures and drug supply, depression in epilepsy is especially easy to miss, which is exactly where a short, epilepsy-specific screen earns its place. Detection matters all the more because access to structured psychological therapy is limited outside metropolitan and tertiary centres, so recognising the disorder and initiating treatment often falls to the treating physician. Stigma compounds the picture: the social consequences of epilepsy for employment and marriage feed low mood directly, and a patient may present the depression as a somatic complaint rather than name it. The iatrogenic contribution of older, widely used antiepileptic drugs to mood is real but belongs to the drug-attributable topic of this chapter and is only flagged here.

22.8 Pulling the picture together

One disorder, higher than you think, stranger than you expect, and quiet unless you ask. The topic resolves into a short chain that is easy to reproduce under pressure. Depression is the commonest interictal psychiatric disorder in epilepsy and the strongest single predictor of poor quality of life, so it deserves the clinical priority that instinct gives to the seizures. Its prevalence appears to scatter wildly only until the scatter is decoded as the product of measure, definition and sampling, at which point the ascending gradient from community symptoms to presurgical major depression becomes orderly and the seizure-control gradient becomes intelligible. The mood disorder is unipolar, not bipolar, and it is frequently atypical and pleomorphic rather than criterion-complete, which is the fact that both explains why it is missed and warns against dismissing it when full criteria are absent. And because it is missed, the clinical answer is to screen actively, with an epilepsy-specific instrument built to see past the somatic noise, handing the confirmed case onward to the treatment and suicide-risk topics that own those questions.

commonest	17 to 21%	24 to 35%	top predictor
INTERICTAL PSYCHIATRIC DISORDER	MAJOR DEPRESSION, EPILEPSY CLINICS	LIFETIME, PRESURGICAL TLE	OF POOR QUALITY OF LIFE

Depression is the commonest and most disabling interictal psychiatric comorbidity of epilepsy, is unipolar rather than bipolar, is frequently atypical and criterion-incomplete rather than textbook major depression, and is missed unless it is actively screened for with an epilepsy-specific instrument such as the Neurological Disorders Depression Inventory for Epilepsy.

23 · Suicide risk in epilepsy and contributing factors

People with epilepsy take their own lives more often than the general population does, and the whole of this topic is an exercise in holding that fact firmly without inflating it. Suicide is one of the genuine hazards of living with epilepsy, and the excess appears to be **intrinsic to epilepsy** itself rather than merely a by-product of the drugs used to treat it: the risk is raised above that of the general population, a burden the disease carries in its own right, though the reported rates vary widely with the population studied. That last distinction is where marks are won and lost, because a separate and quite different suicidality signal attaches to the antiseizure drugs as a class, and it is developed under its own dedicated topic in this chapter. This topic owns the disease-intrinsic risk; the drug-attributable warning is not reopened here, and a reader who leaves with the two collapsed into one has missed the central discrimination the examiner is looking for.

The examinable material divides into three questions that a competent clinician actually asks. The first is how large the excess really is, which turns out to be one of the most contested figures in the whole of epilepsy psychiatry, with reputable estimates ranging from a several-fold increase down to none at all. The second is where suicide sits within the wider excess mortality of epilepsy, which is best read through an all-cause mortality statistic rather than any single suicide figure. The third, and the most clinically useful, is which patients actually carry the risk, and the striking, repeatedly examined answer is that it is features of the person and their history, not features of the epilepsy, that predict it. The source throughout is Lishman's Organic Psychiatry, and the honest posture is to quote the few established figures precisely and to reason carefully around the uncertainty rather than paper over it.

23.1 A real risk, and why it belongs to the disease

Suicide is a bona fide cause of premature death in epilepsy, and the excess is the disease's own. Before arguing about the size of the risk it is worth being clear that the direction is not in doubt. People with epilepsy are, taken as a whole, more likely to die by suicide than their peers without epilepsy, and this holds even when the contribution of any particular drug is set aside, which is what it means to call the risk intrinsic. Several strands plausibly converge to produce it. Epilepsy carries a heavy load of psychiatric comorbidity, with depression the commonest interictal disorder and a powerful driver of hopelessness, and that mood picture is developed in

its own dedicated topic in this chapter. The temporolimbic disturbance that underlies so much of the psychiatry of epilepsy may itself lower the threshold for despair and impulsive self-harm. Layered onto the neurobiology is a formidable psychosocial burden: the unpredictability of the fits, the loss of autonomy and driving, the constraints on employment and marriage, and the stigma that still attaches to the diagnosis, all of which feed a reactive demoralisation that no anticonvulsant reaches. The clinical consequence is that suicide risk in epilepsy cannot be reasoned about as though it were purely reactive or purely organic; it is both, and the assessment has to look at the person as well as the seizures.

RULE **Suicide risk is genuinely elevated in epilepsy and the elevation is intrinsic to the disease, but the size of the excess is contested and the risk is carried by the patient's psychiatric history, not by any feature of the epilepsy itself.** Every clinical move that follows descends from that one sentence: assess the person, not the seizure type, and keep the disease-intrinsic risk separate from the drug-attributable signal owned elsewhere.

23.2 How large the excess is - the figures that will not agree

The reported magnitude swings from a large multiple to nothing, and the swing is information, not noise. No single relative-risk figure for suicide in epilepsy is defensible, and a candidate who offers one confidently is usually offering a wrong one. The disagreement in the literature is real and instructive. An early and much-cited clinic-based review put the excess at **roughly fivefold**, rising to as much as 25-fold in the most highly selected samples such as specialist referral and presurgical populations. Set against that, two **population-based** mortality studies found no excess suicide risk at all. The two sets of numbers are not so much contradictory as answers to different questions. Estimates built from selected clinic and referral cohorts concentrate the sickest, most refractory and most psychiatrically comorbid patients, and in that enriched group the suicide rate is genuinely high; population-based samples dilute those patients back into the whole epilepsy population and the signal shrinks, sometimes to invisibility. The accompanying table sets the two poles side by side. The lesson to carry into the examination is comparative and mechanistic: name the range, explain that selection drives it, and do not commit to a single multiplier as though the field had settled on one.

SOURCE OF THE ESTIMATE	REPORTED SUICIDE RISK VERSUS THE GENERAL POPULATION
Early clinic-based and selected-sample reviews	about fivefold, up to 25-fold in the most highly selected samples
Two population-based mortality studies	no increase over the general population

The two poles of the suicide-risk literature in epilepsy. The discordance is driven by how the sample was assembled, not by genuine disagreement about the disease: selected clinic cohorts concentrate the highest-risk patients, while population-based studies dilute them.

23.3 Suicide within the excess mortality of epilepsy

Read the burden through the all-cause mortality statistic, and keep suicide as one contributor among several. A cleaner way to situate suicide is to step back to the overall mortality of epilepsy, which is unambiguously raised. Expressed as a **standardised mortality**

ratio, the death rate of people with epilepsy runs at about **2 to 3** times that expected in the general population. Crucially, that figure is all-cause, and suicide is only one of several contributors to the excess mortality, alongside the underlying cause of the epilepsy, accidental injury, status epilepticus and sudden unexpected death in epilepsy. Holding the statistic in that form does two useful things. It anchors the reality of premature death in epilepsy in a robust, population-level number rather than the wobbly suicide multiples of the previous section, and it forces the discipline of attribution: the raised overall mortality is not a suicide figure, and the raised suicide risk is not the whole of the raised mortality. Status epilepticus and sudden unexpected death in epilepsy are substantial causes in their own right and are not developed here. The examinable point is proportion and provenance: suicide is a real and consistent slice of the excess mortality, but it is a slice, and the standardised mortality ratio is the frame that keeps the slice honest.

TRAP **Quoting the standardised mortality ratio of 2 to 3 as if it were the suicide risk.** It is not. That ratio is the all-cause excess mortality of epilepsy, of which suicide is one contributor among several. Candidates who have memorised the number attach it to whatever the question asks about, and here they misread an all-cause statistic as a suicide-specific one. The two are different measures and the examiner is testing whether you know which is which.

23.4 A history of attempted suicide is commoner too

Non-fatal self-harm shadows the completed-suicide picture, with the same caveats about how it was counted. The excess is not confined to completed suicide. A past history of attempted suicide is about four times more common in people with epilepsy than in the general population, which matters clinically because a previous attempt is one of the most powerful predictors of an eventual completed suicide in any population and, as the next section shows, in epilepsy specifically. The figure has to be handled with the same care as the completed-suicide multiples, because it comes from survey data with real selection-bias limitations: only about one-third of the questionnaires were returned, and non-responders in a survey of this kind may differ systematically from responders in ways that pull the estimate up or down. The direction is nonetheless the clinically important message, and it is consistent with everything else in this topic. A raised prevalence of prior attempts tells the clinician that self-harm is part of the lived texture of epilepsy for a meaningful minority, and it flags the single item of history that most deserves to be asked about at every review, because it converts an abstract population risk into an individual one sitting in front of you.

23.5 The factors that actually carry the risk

The risk lives in the patient's mind and history, and its strongest markers are the ones that predict suicide anywhere. This is the heart of the topic and the part most worth memorising, because it converts a contested epidemiology into an actionable risk assessment. In a case-control study of suicide in epilepsy, a **psychiatric history** conferred an **about ninefold** increase in suicide risk, and when the analysis was reduced to essentials, a psychiatric history and a **previous suicide attempt** emerged as, by a wide margin, the two most important risk factors for completed suicide. A second, more epilepsy-specific marker sits alongside them: **early-onset**

epilepsy markedly raises the risk, with onset before 18 years, set against onset after 29 years, carrying a **relative risk of 16**, though with a wide confidence interval (4.4 to 58.3) that reflects how few events anchor the estimate; onset in adolescence appears particularly important. The accompanying table lays the three out together. What unifies them is that the first two are exactly the factors that predict suicide in the general population, which is a reassuring internal consistency, while the third is the one genuinely epilepsy-linked contributor and a reminder that a diagnosis made in a developing adolescent brain, at a formative point for identity and independence, carries a weight the same diagnosis made in mid-life does not.

CONTRIBUTING FACTOR	EFFECT ON SUICIDE RISK
A psychiatric history	about a ninefold increase
Onset of epilepsy before 18 years (adolescence particularly)	relative risk of 16, with a wide confidence interval
A previous suicide attempt	one of the two most important factors, with psychiatric history

The contributing factors that carry the elevated suicide risk in epilepsy. Two of the three are the classic general-population predictors of suicide; the third, early age at onset, is the epilepsy-specific contributor. The variables that do not predict suicide are set out in the following section.

MNEMONIC

PAY attention to the person, not the fits

The three factors that raise suicide risk in epilepsy are a **P**sychiatric history, a prior **A**ttempt, and a **Y**oung age at onset. All three are features of the patient and their history, which is precisely why the epilepsy's own characteristics add nothing to the prediction. If you remember to look at the person rather than the seizure chart, you have the risk assessment.

GOOD TO REMEMBER

At every review of a patient with epilepsy, ask directly about past and present psychiatric illness and about any previous suicide attempt, because those two questions, not the seizure type or the drug chart, are what actually stratify suicide risk. Documenting that you asked, and asking again when mood or circumstances change, is the single highest-yield thing this topic asks you to do on a Monday morning.

23.6 The variables that do not predict it

An examinable set of negatives: the epilepsy's own features are largely silent on suicide risk. Just as important as knowing what raises the risk is knowing what does not, because intuition and older theorising both push clinicians to over-read the seizure characteristics. The same case-control work found no consistent relationship between suicide and the type of epilepsy, its localisation or lateralisation, the presence of a neurological deficit, or any specific antiepileptic drug. A trend towards more frequent seizures and polytherapy fell short of statistical significance, so neither can be asserted as an established risk factor. The message is disciplined and counterintuitive. A left temporal focus, a refractory course, a heavy drug burden - the features a

clinician instinctively treats as ominous - do not, on this evidence, mark out the patients who will die by suicide. The seizure-type and localisation detail belongs to the classification and temporal-lobe topics of this chapter and is named here only to be set aside as non-contributory. The list below is the compact form worth reproducing under pressure.

- The type of epilepsy (focal versus generalised, and syndrome).
- The localisation or lateralisation of the focus.
- The presence of a neurological deficit.
- Any specific antiepileptic drug taken for seizure control.
- Seizure frequency and polytherapy, where a trend was seen but did not reach significance.

PITFALL

The clinical error this section guards against is assuming that good seizure control equals safety. Because suicide risk tracks the psychiatric history and prior attempts rather than the seizures, a patient rendered seizure-free can remain at substantial risk, and a clinic congratulating itself on seizure freedom may quietly stop asking about mood. Seizure control is necessary for a good outcome but is not sufficient for safety, and the risk assessment must run independently of the seizure chart.

23.7 Turning the risk picture into clinical action

Everything above converges on a single practical instruction: assess and treat the person, not the seizure type. The management of suicide risk in epilepsy is not a distinct pharmacology but a disciplined application of general suicide-prevention practice to a higher-risk group, framed by the epilepsy standards of the International League Against Epilepsy and NICE and by the epilepsy-psychiatry management literature synthesised in Lishman's Organic Psychiatry, which together set optimisation of seizure care and active attention to mental health as twin obligations. The **LOCUS** of care is overwhelmingly outpatient, woven into routine epilepsy or psychiatric follow-up, and escalates to the emergency department or inpatient admission only when active suicidal ideation with intent or a recent attempt makes immediate safety the priority, at which point management becomes that of any acute suicidality. The focus shifts by phase: the acute task is to detect and contain active ideation and to secure safety; the short and mid-term task is to treat the comorbid psychiatric disorder that is driving the risk, most often depression, and to keep the seizures well controlled; and the long-term task is to sustain vigilance across a chronic, relapsing illness rather than assess risk once and file it. The modus spans every available mode without changing the antiseizure regimen for the sake of the mood. The drug mode is the vigorous treatment of the underlying psychiatric illness, with the specific question of which antidepressant to choose against the seizure threshold handed to its own dedicated topic in this chapter. The psychological mode is structured risk assessment and collaborative safety planning. The social mode is family involvement, means restriction and reliable continuity of follow-up. The elements below are the workable core.

- Ask every patient about psychiatric history and any previous suicide attempt, and re-ask when circumstances change.
- Treat comorbid depression and other psychiatric illness actively rather than dismissing low mood as understandable.
- Restrict access to means and engage the family as partners in monitoring between appointments.
- Keep seizure control optimised while remembering it does not, by itself, reduce suicide risk.

IN INDIA

The risk assessment matters all the more where follow-up is hurried and access to structured psychiatric care is thin outside tertiary centres, so the treating physician often has to carry the mental-health task alone. Means restriction has a specific local shape: ready access to pesticides and other agricultural poisons in rural settings makes counselling the family about safe storage a concrete and potentially life-saving intervention rather than an abstraction. Stigma around both epilepsy and suicide runs deep and makes families slow to volunteer a change in mood or a past attempt, so the questions have to be asked directly, in the patient's own language, and an accessible contact route, often a phone or messaging line to the treating team, is a more realistic safety net than a distant review date.

23.8 Pulling the picture together

A real excess, an unsettled size, and a risk that sits in the person. The topic resolves into a short chain that is easy to reproduce under pressure. Suicide risk in epilepsy is genuinely elevated and the elevation is intrinsic to the disease, but it must be kept apart from the separate drug-attributable suicidality signal taught elsewhere in this chapter. The size of the excess is one of the field's most contested figures, running from a large multiple in selected clinic samples to nothing at all in population-based studies, and the discordance is explained by selection rather than by any real disagreement about the disease. The overall mortality of epilepsy is unambiguously raised, and suicide is one consistent contributor to that excess among several, best read through an all-cause statistic that must not be misquoted as the suicide risk itself. And the contributing factors that actually predict completed suicide are a psychiatric history and a previous suicide attempt, joined by early age at onset, while the type of epilepsy, its localisation, the neurological deficit and the specific drug do not predict it. Assess the person, treat the comorbidity, and never let seizure freedom stand in for safety.

Elevated	**Contested**	**Psychiatric history**	**Not the epilepsy**
AND INTRINSIC TO THE DISEASE	CLINIC-HIGH, POPULATION-NULL	THE DOMINANT RISK FACTOR	TYPE, FOCUS AND DRUG DO NOT PREDICT

Suicide risk in epilepsy is genuinely elevated and intrinsic to the disease, though its size is contested from a several-fold excess down to none; the risk is carried by a psychiatric history and a previous suicide attempt, joined by early age at onset, and not by the epilepsy type, localisation, neurological deficit or specific drug, and it must never be conflated with the all-cause standardised mortality ratio or the separate drug-attributable suicidality signal.

24 · Anxiety disorders in epilepsy

Anxiety is the quiet comorbidity of epilepsy. The anxiety disorders in epilepsy divide into a peri-ictal group, in which seizure-related fear is the prototype, and an interictal group. The mood-in-epilepsy literature is dominated by depression, and the psychoses draw the eye because they are dramatic, but anxiety runs alongside both and is, on the best available evidence, at least their equal in frequency. It is named far less often than it is felt, partly because its symptoms are so easily written off as an understandable reaction to an unpredictable, stigmatising illness, and partly because the whole field learned to screen for depression first. The examinable content sits in four places, and they map cleanly onto four things a good clinician actually does: knowing how common interictal anxiety really is, telling a genuine anxiety disorder apart from an ictal fear that is really a seizure, recognising the epilepsy-specific entity of seizure phobia, and treating the picture without entrenching the benzodiazepine misuse that so often complicates it.

These presentations are best sorted, as the rest of the peri-ictal material in this chapter is, by their relationship in time to the seizure. Anxiety arising between seizures is a true psychiatric comorbidity and is treated as such; anxiety that is the seizure, the sudden fear or panic that begins and ends with the discharge, is a neurological event wearing a psychiatric mask and belongs to a separate topic. Almost every mistake in this area comes from letting those two collapse into each other, so the discrimination is the spine on which the rest of the answer hangs.

24.1 How common it is, and why it is missed

At least as common as depression, and consistently under-recognised. Lishman's Organic Psychiatry, drawing together community and clinic-based studies, concludes that in the interictal state anxiety and depression are probably about equally prevalent, with several series finding anxiety symptoms to be a little more common than depressive ones. That comparison is the one worth carrying, because it corrects the instinct to treat anxiety as the lesser problem. The absolute figure is deliberately left qualitative: reported rates swing so widely with the population sampled, from community cohorts through neurology clinics to surgical series, that a solitary percentage without its denominator is a weak and probably wrong answer. The safe form of the fact is comparative, not numeric. Under-recognition has identifiable causes. Anxious symptoms are readily attributed to the epilepsy itself or to the antiepileptic drugs, or dismissed as the sensible worry of someone who never knows when the next fit will come, and so are normalised out of the picture. The comorbidity is also crowded out by the two conditions it accompanies, a depression that is actively screened for and a psychosis that demands attention, each taught in its own right elsewhere in this chapter. The cost of missing it is not trivial: untreated anxiety erodes quality of life at least as much as seizure frequency does, inflates unscheduled care, and

compounds both depression and the elevated suicide risk that epilepsy carries, each developed in its own dedicated topic. Anxiety is not a soft add-on to the mood picture but a hard, treatable target in its own right.

RULE **Anxiety is at least as frequent as depression in epilepsy, and it is treatable.** Screen for it deliberately rather than waiting for it to be volunteered, because it is routinely normalised as an understandable reaction and so goes unnamed and untreated. The comparative fact, that it equals or slightly exceeds depression, is the examinable one; a single point-prevalence figure is not.

24.2 Interictal anxiety versus ictal fear: the cardinal distinction

Is this an anxiety disorder, or is it a seizure? The first decision about any anxious presentation in epilepsy is where it sits on the seizure timeline, because two very different things look alike. An **interictal anxiety** disorder is a comorbidity arising between seizures and resembling anxiety in anyone else. **Ictal fear or ictal panic**, by contrast, is a seizure phenomenon and not an interictal anxiety disorder at all; it is a brief, stereotyped, self-limiting affective event that is the seizure, and its neurobiology and management are set out in their own topic in this chapter, so it is named here only as the thing to exclude. The distinction is not hair-splitting: it dictates the entire workup. A true anxiety disorder is managed as a comorbidity, whereas a paroxysm of ictal fear is a call to optimise seizure control, and an antidepressant will not touch it. Several bedside features do the separating work. Ictal events are brief, sudden, stereotyped from episode to episode, often accompanied by other seizure phenomena such as automatisms or altered awareness and followed by postictal features; interictal anxiety is sustained, situationally shaped, elaborated with worry and anticipation, and shows none of that paroxysmal stereotypy. The accompanying figure places the common anxiety-like presentations against the seizure to make the sorting explicit.

PRESENTATION	TIMING TO THE SEIZURE	WHAT IT ACTUALLY IS	WHERE IT IS TAUGHT
Interictal anxiety disorder	between seizures	a genuine anxiety disorder, like anxiety in the general population	here
Ictal fear or ictal panic	during the seizure	a seizure phenomenon, not an anxiety disorder	its own topic in this chapter
Seizure phobia	between seizures	fear of having a seizure, out of proportion to the real risk	here
Anticipatory anxiety	between seizures	understandable worry about an unpredictable illness, on a continuum with the above	here

Anxiety-like presentations in epilepsy sorted by timing to the seizure; three of the four are interictal comorbidities managed here, while ictal fear is a seizure named only as the differential to exclude.

TRAP **Treating a seizure as an anxiety disorder, or an anxiety disorder as a seizure.** Candidates lose the mark by starting an antidepressant for paroxysmal ictal fear, or by escalating antiepileptic drugs for a sustained interictal anxiety disorder. The tell is stereotypy and brevity: an event that is identical each time, lasts seconds to a minute or two, and carries other seizure features is ictal, not an anxiety disorder, however frightening it feels.

24.3 Which anxiety disorders predominate

Agoraphobia, generalised anxiety and social phobia lead; obsessive-compulsive disorder does not. Broken down by diagnosis, a consistent pattern emerges. Lishman, following the work of **Jones and colleagues (2005)**, reports that the most frequent categories are **agoraphobia**, **generalised anxiety disorder** and **social phobia**. The clustering makes clinical sense from the patient's side. Agoraphobia and social avoidance are the natural sequelae of an unpredictable event that may strike in public, be witnessed, and be humiliating or dangerous, so the person narrows their world to places felt to be safe; generalised anxiety is the diffuse, persistent apprehension of living with an illness that gives little warning. Against that grain sits a negative fact that is examinable precisely because it is counter-intuitive: **obsessive-compulsive disorder** is not over-represented in epilepsy. Older teaching linking temporal lobe and limbic pathology to obsessional phenomena might lead a candidate to expect an excess of it, and the discipline of the topic is to resist that expectation and state the negative. The pattern is worth holding whole: the disorders that flourish are those an unpredictable, socially exposed illness would breed, and the one that does not is the one intuition wrongly predicts. Naming the categories also sharpens management, since agoraphobic and social avoidance respond particularly well to the exposure-based psychological work discussed below, whereas diffuse generalised anxiety is the picture in which a drug is more often needed.

MNEMONIC

A-G-S, not O

The interictal anxiety disorders that cluster in epilepsy are **A**goraphobia, **G**eneralised anxiety disorder and **S**ocial phobia. The letter deliberately left out is **O** for obsessive-compulsive disorder, which is the one that is not over-represented. Each element is a fact the topic actually states, and the omission is as examinable as the inclusions.

24.4 Seizure phobia: the epilepsy-specific syndrome

When the fear of the fit becomes more handicapping than the fit. The one entity in this area with no counterpart in general psychiatry is seizure phobia, and it is the part most likely to be tested. It is best defined as a fear of having a seizure that is out of all proportion to the objective risk and that disables the patient more than the seizures themselves do. Phenomenologically it resembles agoraphobia: the dreaded outcome is a seizure in a place where it would cause injury or public embarrassment, and the response is to avoid such situations and lean on safety behaviours. That structure explains its self-perpetuating quality. Every avoided outing that passes without a fit is misread as proof that avoidance works, so the avoidance is reinforced and the world shrinks further, exactly as in any phobic disorder. The most instructive feature, and the one showing the

fear has become autonomous rather than a rational appraisal of danger, is that seizure phobia can be at its most striking after successful epilepsy surgery, in patients who are seizure-free yet remain housebound by dread of a fit the operation has abolished. The fear has outlived its object. That is the clinching argument for treating the anxiety directly rather than assuming it will melt away once seizures are controlled; it will not, because it has detached from them. Seizure phobia should therefore be actively asked about, since patients rarely volunteer a fear they regard as sensible, and treated as a phobic disorder in its own right rather than an inevitable shadow of the epilepsy.

24.5 Where the anxiety comes from

Reactive, organic and iatrogenic strands, usually braided together. Anxiety in epilepsy is rarely traceable to a single cause, and thinking in strands helps both formulation and treatment. The reactive strand is the most visible: living with an illness that strikes without warning, carries stigma, can injure, and constrains driving, employment and independence generates a realistic apprehension that can escalate into a disorder. The organic strand reflects the same limbic and cortical disturbance that underlies the epilepsy, so that in some patients the anxiety is a direct expression of the brain's dysfunction rather than a response to it, on the neurobiological substrate this chapter develops in its dedicated mechanism topic. The iatrogenic strand is treatment itself, since antiepileptic drugs have their own behavioural and psychiatric adverse effects, the detail of which is owned by a separate topic in this chapter. In practice the strands coexist, and the value of separating them is that each points to a different lever.

- The unpredictability and perceived dangerousness of seizures, which drives anticipatory anxiety and phobic avoidance.
- Stigma and the social consequences of a witnessed fit, which feed social phobia and agoraphobic constriction.
- The direct neurobiological disturbance shared with the seizure disorder, which can generate anxiety autonomously.
- The psychiatric adverse effects of antiepileptic drugs, an iatrogenic contribution that is always worth reviewing.

Separating the strands is not a formality: a reactive, phobic picture calls first for psychological work, a predominantly iatrogenic one for a hard look at the drug regimen, and an organic anxiety yoked to poor seizure control for better seizure control above all. Most patients need attention to more than one strand, which is why the management below spans several modes at once.

24.6 The benzodiazepine complication

The picture is frequently muddied by benzodiazepine dependence. A specific and common complication deserves its own place because it changes the whole management calculus: the anxiety of epilepsy is often layered over, and worsened by, **benzodiazepine dependence**. Two routes lead there, both easy to walk into. In the first, rescue benzodiazepines legitimately prescribed for seizure clusters come to be taken for supposed prodromes, so a drug meant for an emergency is used pre-emptively and repeatedly for a warning that may be nothing more than

anticipatory anxiety. In the second, the patient reaches for the same drugs as self-medication for the anxiety itself, finding short-term relief that narrows into dependence. Benzodiazepine use and anxiety then drive each other: inter-dose and withdrawal anxiety mimics and amplifies the underlying disorder, escalating use looks like worsening illness, and the sedation and tolerance that follow do nothing for the phobic core. Recognising the loop matters because the correct move, a graded withdrawal alongside proper treatment of the anxiety, runs against the patient's instinct and often against the path of least resistance in a busy clinic. An anxious person with epilepsy who is escalating benzodiazepines is not usually one who needs more of them.

PITFALL

The reflex response to a distressed, anxious patient with epilepsy is to reach for a benzodiazepine, or to intensify the antiepileptic regimen, when the actual problem is a phobic anxiety disorder or seizure phobia that needs psychological treatment. Each reflex entrenches the wrong thing: benzodiazepines breed dependence and antiepileptic escalation adds adverse effects, while the fear that is doing the disabling goes untreated. Name the anxiety disorder before you medicate the seizure.

24.7 Management

Treat it as anxiety, and treat the epilepsy as epilepsy. The governing principle is that anxiety in epilepsy is managed broadly as anxiety is in general psychiatry, with the epilepsy-specific overlays of seizure phobia, the ictal-fear differential and the benzodiazepine trap on top. The work is framed by the epilepsy standards of the International League Against Epilepsy and NICE, which make optimised seizure control the foundation, together with the epilepsy-psychiatry literature synthesised in Lishman's Organic Psychiatry and Kaplan and Sadock's Comprehensive Textbook of Psychiatry for the anxiety itself. The **LOCUS** of care is overwhelmingly outpatient, since interictal anxiety and seizure phobia are chronic, ambulatory problems; the exceptions that pull care into a more intensive setting are a difficult planned benzodiazepine withdrawal and the rare presentation severe enough to threaten safety or function. The **focus** shifts by phase: the acute task is to contain distress and, critically, avoid entrenching benzodiazepine use; the short-to-mid-term task is to treat the specific anxiety disorder or seizure phobia to remission; and the long-term task is relapse prevention through durable seizure control, a benzodiazepine-sparing regimen and attention to driving, work and stigma.

The **modus** spans every available mode. Psychological treatment leads, and it is where seizure phobia in particular shines, since it responds well to **cognitive-behaviour therapy**, the approach established for it by **Newsom-Davis and colleagues (1998)**; the therapy challenges the catastrophic cognitions and schema the patient holds about their epilepsy and dismantles the safety and avoidant behaviours that keep the phobia alive. Pharmacologically, anxiety in epilepsy is treated as elsewhere, with an SSRI as first-line; where an SSRI has not delivered, pregabalin is a reasonable alternative, having shown an antianxiety effect in epilepsy patients with generalised anxiety disorder, though on the limited evidence of a small open-label study rather than anything definitive. The choice of any antidepressant against its effect on the seizure threshold is a

decision handled in its own dedicated topic in this chapter and is not reopened here. The procedural and social modes complete the plan: optimising antiepileptic treatment as the organic lever, withdrawing benzodiazepines where dependence has taken hold, and addressing the driving, employment and family realities that sustain the reactive anxiety. Running through all of it is the caution that the drugs offering the quickest relief are the ones most likely to make the long-term picture worse.

IN INDIA

Two features sharpen this topic in Indian practice. First, benzodiazepines, and clonazepam in particular, are widely and easily accessible, so the slide from rescue or as-needed use into dependence is an everyday risk and the benzodiazepine-sparing discipline above is not theoretical. Second, trained cognitive-behaviour therapists are scarce outside metropolitan and tertiary centres, a real problem precisely because seizure phobia and the phobic-avoidant disorders respond so well to that treatment; the practical answer is often SSRI treatment supported by whatever structured psychological input can be mustered, alongside direct work on the stigma around epilepsy that, given employment and marriage prospects, does much of the driving here.

GOOD TO REMEMBER

Ask every patient with epilepsy two questions the notes rarely capture: how much do you fear the next seizure, and how often do you take a benzodiazepine for the fear. The first uncovers seizure phobia and the phobic-avoidant disorders that dominate the picture; the second uncovers the dependence that quietly worsens it. Both are missable, both are common, and both are treatable.

24.8 Pulling the picture together

One axis, three real disorders, one impostor. The whole area resolves if the anxious presentation is first placed against the seizure. Everything interictal, the generalised, agoraphobic and social anxiety that cluster in epilepsy, the seizure phobia that can outlast the fits it fears, and the ordinary anticipatory dread of an unpredictable illness, is a genuine, treatable psychiatric problem. The one impostor is ictal fear, the paroxysm that is the seizure rather than a reaction to it, named here only to be sent to its own topic. Hold that distinction, remember that anxiety at least matches depression in frequency and is routinely under-treated, recall that obsessive-compulsive disorder is the exception that does not cluster, and keep one hand on the benzodiazepines, and the topic answers itself.

≥ depression	Agora / GAD / Social	OCD not raised	CBT + SSRI
INTERICTAL ANXIETY FREQUENCY	COMMONEST ANXIETY DISORDERS	NOT OVER-REPRESENTED	FIRST-LINE TREATMENT

Interictal anxiety is at least as common as depression in epilepsy and is genuinely treatable; agoraphobia, generalised anxiety and social phobia predominate while obsessive-compulsive disorder does not, the epilepsy-specific syndrome is seizure phobia (fear of the fit that outstrips the fit and can persist after successful surgery), and the two governing cautions are to separate a true anxiety disorder from ictal fear and to keep the anxiety from becoming benzodiazepine dependence.

25 · Antidepressant choice and seizure-threshold safety in epilepsy

Depression is the commonest psychiatric comorbidity of epilepsy, it is treatable, and it is repeatedly left untreated for the wrong reason. The reason is an old anxiety that antidepressants lower the seizure threshold and will therefore make an already seizure-prone brain worse. That anxiety is not baseless, because some antidepressants genuinely do lower the seizure threshold, but it has been inflated far beyond the evidence, and the price of the inflation is a patient with disabling, sometimes lethal, depression who is denied an effective drug. The marks in this topic are not for reciting a list of dangerous agents. They are for holding two facts at once: that at ordinary treatment doses the seizure risk of a well-chosen antidepressant is very small, and that the choice of agent nonetheless matters, because a few specific drugs carry a real and avoidable liability.

This is a prescribing topic, so it is examined as a decision rather than as a description. The prevalence, phenomenology and screening of depression in epilepsy are developed in their own topic, as is the elevated suicide risk that makes under-treatment dangerous; here the question is what to prescribe once depression has been identified, how to start it, when a physical treatment is preferable, and what to check before the first prescription. The mechanism by which drugs lower the seizure threshold is itself taught separately, in the dedicated topic on drug-induced lowering of the seizure threshold; this topic applies that principle to pick an agent rather than re-deriving it.

25.1 The first principle: treat the depression

The proconvulsant fear has been overstated, and the untreated illness is the greater danger. The historical reluctance to give antidepressants in epilepsy rested on case reports of seizures and on the true observation that these drugs can lower the seizure threshold, but when the actual incidence is examined the picture is reassuring. Lishman's Organic Psychiatry, drawing the figure from non-epileptic populations at therapeutic doses, puts the seizure incidence at **well below 1%**, a rate that has to be set against the substantial morbidity, and the mortality, of an epilepsy patient whose depression goes untreated. The clinical inference is direct: clinicians should not hesitate to prescribe an antidepressant when depression is present and otherwise warrants treatment, because the seizure hazard at ordinary doses is small and the illness being treated is not. Kaplan and Sadock's Comprehensive Textbook of Psychiatry frames the same reassurance mechanistically. The more recently introduced antidepressants are only marginally proconvulsant, and at therapeutic doses the tricyclics, the newer agents and the SSRIs do not usually promote seizures at all; it is **overdose**, not the therapeutic dose, that is the main seizure

trigger. That single distinction dissolves most of the historical worry, because it locates the danger in intoxication rather than in treatment, and it reframes the prescribing task as one of choosing sensibly and starting cautiously rather than of avoiding the class outright.

RULE **Do not withhold an antidepressant from a depressed patient because they have epilepsy.** At therapeutic doses the seizure risk of a well-chosen agent is very low, and untreated depression carries its own serious morbidity and mortality. The live decision is which drug and how to start it, not whether to treat.

25.2 Choosing the agent: SSRIs first

Comparable efficacy means the decision is driven by safety and tolerability. Across the antidepressant classes the efficacies for depression are broadly similar, so in a patient with epilepsy the choice is settled less by potency than by which agent is least likely to provoke a seizure and best tolerated. On that logic the **SSRIs** are the antidepressants of choice and stand first-line, because they combine the usual antidepressant efficacy with a relatively low seizure liability. Beyond the SSRIs a wider group of newer agents is regarded as carrying a comparably low risk and is drawn on where an SSRI is unsuitable: the **lower-seizure-risk antidepressants** named by Lishman's Organic Psychiatry are citalopram, escitalopram, fluoxetine, paroxetine and sertraline among the SSRIs, with mirtazapine, moclobemide, reboxetine, trazodone, venlafaxine and nefazodone beyond them. Because efficacy does not separate them, the practical selection within this group turns on a small number of considerations.

- Seizure-threshold safety, the property that removes the highest-liability agents from consideration altogether.
- The side-effect fit for the individual, since tolerability is what efficacy-equivalence leaves to decide.
- The interaction profile with the patient's antiepileptic drugs.
- Safety in overdose, which weighs more heavily here because depression in epilepsy carries a raised suicide risk.

TIER	CLASS	SEIZURE-THRESHOLD LIABILITY	PRACTICAL CAVEAT
First-line	SSRIs and the wider lower-risk group	lowest	still start low and titrate slowly
Second-line	tricyclics	higher, dose-related	higher risk but not contraindicated
Usually avoided	the highest-liability agents	greatest threshold lowering	bupropion is formally contraindicated in seizure disorders

The antidepressant selection ladder in epilepsy: efficacy is broadly comparable across the tiers, so the ordering is by seizure-threshold liability and practical caveat rather than by antidepressant potency.

25.3 Tricyclics: second-line, not forbidden

A higher, dose-related risk that still leaves them usable. The **tricyclic antidepressants** sit second-line in epilepsy. They carry a higher risk of seizure exacerbation than the SSRIs, and that risk is dose-related, but, and this is the examinable subtlety, they are not contraindicated. A tricyclic remains a legitimate choice when the depression has not responded to safer first-line agents, provided the dose is kept modest and raised slowly. The group named in this second-line position includes amitriptyline, clomipramine, dothiepin, imipramine, nortriptyline, trimipramine and lofepramine. Clomipramine deserves a separate flag, because it is the tricyclic with the greatest proconvulsant liability, and for that reason it appears not only in this second-line list but also among the agents that are usually avoided outright; the same molecule can be both a permissible second-line tricyclic in general use and a drug to steer around specifically in epilepsy. The general lesson of the tricyclic tier is that "second-line" is a real and useful category sitting between "preferred" and "avoid": it says that these drugs are more hazardous than the SSRIs and should not be a first reach, without pretending that they are unusable in a patient who genuinely needs them. Reading the tier correctly, as higher risk, dose-related, and still available, is worth more than memorising the roster, because it is the reasoning an examiner is testing.

25.4 The agents to avoid, and bupropion in particular

A short list of genuine offenders, headed by one that is contraindicated. A small number of antidepressants lower the seizure threshold enough that they are usually avoided in epilepsy altogether: **amoxapine, bupropion, clomipramine and maprotiline**. These are the agents whose threshold-lowering is greatest, and in a patient who already has seizures the sensible default is simply not to use them when an equally effective, safer alternative exists. Bupropion is the one to know in most detail, because its seizure liability is both dose-dependent and, at the top of that list, a formal contraindication in a seizure disorder. Seizures have been reported in about **0.4%** of bupropion-treated patients, and because the effect is dose-related the risk is minimised by using the sustained-release form, by titrating slowly from a low start of **100 mg** daily, and by respecting the sustained-release maximum of **400 mg** daily. In a patient with epilepsy, however, those mitigations are beside the point: the drug is contraindicated in seizure disorders, and in eating disorders where the seizure risk is likewise raised, so the correct action is to choose something else rather than to prescribe it carefully. The principle that unites the four is the one applied throughout this topic, to pick against seizure-threshold liability, taken to its logical end for the agents where that liability is highest.

MNEMONIC

ABCM - the four to avoid

The antidepressants usually avoided in epilepsy, by first letter: **A**moxapine, **B**upropion, **C**lomipramine, **M**aprotiline. Bupropion is the one that is not merely avoided but contraindicated in a seizure disorder. Each name is an agent this topic actually places on the avoid list.

PITFALL

Bupropion reaches patients by routes that bypass the epilepsy history. It is chosen as an activating antidepressant when fatigue or sexual side effects dominate, and it is prescribed for smoking cessation by clinicians who are not thinking about seizures at all. In a person with epilepsy either route is a mistake, because the drug is contraindicated in seizure disorders. Screen for a seizure history before starting bupropion for any indication, not only when treating mood.

25.5 How to start: low dose, slow titration

Because the proconvulsant effect is dose-related, even the safe agents are started cautiously. Within the preferred group the way an antidepressant is introduced still matters, and the governing fact is that the proconvulsant effect of these drugs is dose-related. Lishman's Organic Psychiatry translates that into a simple prescribing rule, to **start low and go slow**, beginning at a low dose and up-titrating slowly, so that the plasma level, and with it any threshold-lowering effect, rises gradually rather than in a jump the brain has to absorb at once. The rule applies across the classes, not only to the second-line tricyclics, because dose-dependence is a general property; it is simply most consequential where the baseline liability is highest. The corollary concerns overdose, and it is the point at which the prescribing decision loops back to the reason for treating in the first place. Therapeutic doses of the tricyclics, the newer agents and the SSRIs do not usually promote seizures, but overdose does, and depression in epilepsy carries a raised risk of self-harm. That combination is a specific argument for the safer agents and for attention to the quantity dispensed, because the very patient in whom a therapeutic dose is low-risk may be at risk of taking far more than the therapeutic dose. Starting low and going slow protects against the therapeutic-dose hazard; choosing an agent that is safer in overdose, and limiting supply, addresses the larger one.

TRAP **Writing that antidepressants are contraindicated in epilepsy.** The blanket statement is wrong and loses the mark. Only a few named agents are usually avoided, and only bupropion carries a formal contraindication in seizure disorders; the SSRIs are first-line and even the tricyclics are second-line rather than forbidden. The examiner wants the graded answer, not a class-wide ban.

25.6 ECT for severe depression

An effective and safe option, and paradoxically unproblematic in a seizure disorder. When depression in epilepsy is severe, whether psychotic, life-threatening through suicidality or refusal to eat and drink, or unresponsive to drugs, **electroconvulsive therapy** is an effective and safe treatment, and it is worth stating plainly because clinicians hesitate over the apparent contradiction of deliberately inducing a seizure in someone whose problem is seizures. Lishman's Organic Psychiatry settles the worry: ECT is effective and safe in this population, and it usually has no discernible effect on the frequency of the patient's own epileptic seizures. Practically, a therapeutic seizure can generally be induced without needing to reduce the patient's antiepileptic drugs first, which removes the main logistical objection to giving it. The reasoning is that the controlled, brief seizure of a treatment session is a different event from the disorder's spontaneous seizures and does not destabilise seizure control in the way intuition fears. This

makes ECT a genuine and sometimes preferable route in exactly the patients where the stakes are highest and where waiting for a slowly titrated drug to work is least acceptable, the severely depressed, the acutely suicidal, and those who have already failed pharmacotherapy. It should therefore be held not as a last resort forced by the failure of everything else, but as a first-rank option for severe depression in epilepsy on its own merits of speed and safety.

GOOD TO REMEMBER

When depression in an epilepsy patient is severe or the patient is acutely suicidal, do not defer electroconvulsive therapy on the assumption that it will worsen the epilepsy. It is effective, it is safe, and it usually does not change seizure frequency, and a treatment seizure can generally be induced without first cutting back the antiepileptic drugs.

25.7 Check interactions before you prescribe

Few matter in practice, but the check is not optional. Before writing the prescription, look for interactions between the chosen antidepressant and the patient's antiepileptic drugs, and in particular the **enzyme-inducing antiepileptic drugs**, whose induction of hepatic metabolism can lower the plasma concentration of a co-prescribed antidepressant. Lishman's Organic Psychiatry is reassuring about the scale of the problem, since clinically significant interactions are few in practice, but reassurance about frequency is not the same as permission to skip the step, because the interactions that do matter can undo treatment by keeping the antidepressant sub-therapeutic. The detailed pharmacokinetics of these interactions, which antiepileptic induces or inhibits which enzyme and how the psychotropic level is affected, are set out in their own dedicated topic on the pharmacokinetic interactions between antiepileptic drugs and psychotropic agents; they are named here only to place the check in the prescribing sequence, not to be reworked. The practical posture is one of routine vigilance rather than alarm: assume most combinations will be fine, confirm that the specific pair in front of you is one of them, and be ready to use a higher antidepressant dose, or to prefer an agent less affected by induction, when an enzyme-inducing antiepileptic is on board. The step costs little and occasionally prevents a treatment failure that would otherwise be misread as a non-responsive depression.

IN INDIA

The interaction check carries more weight in Indian practice than the Western default implies, because the older enzyme-inducing antiepileptic drugs, phenobarbital above all, and phenytoin and carbamazepine, remain widely used first-line agents on grounds of cost and availability. A depressed epilepsy patient is therefore more likely to be on an enzyme inducer than a patient in a setting where the newer non-inducing antiepileptics predominate, and the co-prescribed antidepressant is correspondingly more likely to be pushed toward sub-therapeutic levels. The SSRIs that lead the first-line choice are, helpfully, cheap and readily available, so the preferred agent is rarely the constraint; the discipline that matters is remembering to check the antidepressant against the inducing antiepileptic, rather than assuming, as one safely might elsewhere, that the background drug is metabolically quiet.

25.8 Putting it together: managing depression in the epilepsy patient

The pieces assemble into a single prescribing pathway. The management of depression in epilepsy is anchored in the epilepsy guidelines of ILAE and NICE for the seizure disorder itself, and in the epilepsy-specific psychiatric-management literature together with the Maudsley Prescribing Guidelines for the antidepressant decision, with textbook prose as the connective cover.

The **LOCUS** of care is usually the outpatient clinic, where most depression in epilepsy is identified and treated; inpatient admission is reserved for severe, psychotic or high-suicide-risk presentations, and the emergency department is where acute suicidality first arrives, the elevated suicide risk of epilepsy being handled as its own topic. The **focus** moves by phase: the acute task is to secure safety and begin an effective antidepressant; the short-to-mid-term task is to titrate that agent to response while watching seizure frequency; and the long-term task is maintenance to prevent relapse and continued optimisation of seizure control, since better-controlled epilepsy is itself better for mood. The **modus** spans every mode of treatment. The pharmacological mode is the spine of this topic: an SSRI first, the wider lower-risk group next, tricyclics second-line where needed, the highest-liability agents avoided and bupropion excluded outright, every drug started low and raised slowly, and interactions with enzyme-inducing antiepileptics checked first. The procedural mode is electroconvulsive therapy for severe or refractory depression, effective and safe without destabilising seizures. The psychological mode, structured psychotherapy for the depression, and the social mode, addressing the disability, stigma and loss of driving or employment that feed low mood in epilepsy, complete the plan; the choice of an antipsychotic against its seizure-threshold effect in epileptic psychosis is the sibling question this pathway sits beside.

below 1% SEIZURE RISK AT THERAPEUTIC DOSE	**0.4%** SEIZURES ON BUPROPION	**400 mg** BUPROPION SR MAXIMUM (CONTRAINDICATED IN EPILEPSY)	**SSRIs** FIRST-LINE IN EPILEPSY

In epilepsy, treat the depression: at therapeutic doses the seizure risk of a well-chosen antidepressant is well below 1%, SSRIs are first-line and tricyclics second-line but not contraindicated, amoxapine, bupropion, clomipramine and maprotiline are usually avoided with bupropion contraindicated outright, every agent is started low and titrated slowly, ECT is effective and safe for severe cases, and interactions with enzyme-inducing antiepileptics are checked before the first prescription.

26 · AED-attributable mood and suicidality (FDA warning) and management

The suicidality caution printed on every antiseizure prescription is a pharmacovigilance story, not a piece of psychopathology. The examinable core of this topic is a single regulatory event and the epidemiology that produced it. In **2008** the United States Food and Drug Administration issued a **class warning** of increased **suicidality**, meaning suicidal ideation or behaviour, for antiepileptic drugs taken as a group. The mark-earning discrimination, and the one

a good answer opens with, is that this drug-attributable signal is a different thing from the suicide risk that is intrinsic to epilepsy itself. The elevated suicide risk of the disease, expressed as its standardised mortality ratio, is a separate and in fact larger question that is worked up under its own topic in this chapter, and the reader must not leave with the two collapsed into one.

The topic tests three things at once: how a drug-safety signal is generated from pooled trials, how relative and absolute risk are read against each other, and how a caution of this kind should change prescribing. It is also a topic where confident memory is dangerous, since the exact regulatory form of the warning and the arithmetic of the risk are both routinely misremembered. The safe approach is to hold the few established facts precisely and reason around them, rather than inflate the signal into something it was never shown to be.

26.1 What the 2008 warning actually says

A class-level caution about emergent suicidality, applied to the whole drug group. The warning is best understood first for what kind of statement it is. It is not a diagnosis, not a claim that antiseizure drugs cause a depressive illness, and not a statement about any single agent. It is a **class effect**: an increase in the frequency of suicidal thoughts and suicidal behaviour observed across antiseizure drugs considered together, applying to patients treated for epilepsy and to those with psychiatric disorders alike, as summarised in Kaplan and Sadock's Comprehensive Textbook of Psychiatry. The endpoint of interest is deliberately behavioural and cognitive rather than diagnostic. It captures new or worsening suicidal ideation and suicidal acts as adverse events emerging on treatment, which is why the clinical response the warning asks for is monitoring rather than a change in diagnosis. Framing it this way also keeps the topic in its lane. The negative affective and behavioural effects that particular agents can produce, and the broader cognitive and behavioural adverse effects of antiseizure drugs, are owned by their own dedicated topics in this chapter and are only named here; what belongs to this topic is the class-wide suicidality signal and the regulatory and clinical machinery built around it. A candidate who starts reciting which individual drug is most depressogenic has drifted off the actual question, which is about a pooled safety signal and what a prescriber does with it.

26.2 The evidence base: a pooled analysis of placebo-controlled trials

The signal came from meta-analysis, not from spontaneous reports. The strength of the warning, and much of its credibility, rests on how it was generated. Rather than resting on case reports or registry noise, the FDA assembled a pooled analysis of **nearly 200 placebo-controlled trials** spanning eleven antiseizure drugs, drawing together tens of thousands of randomised patients, as tabulated in Dulcan's Textbook of Child and Adolescent Psychiatry. Pooling randomised, placebo-controlled data matters because it is the design least vulnerable to the biases that plague drug-safety questions: randomisation balances the baseline suicide risk between arms, the placebo comparator supplies the counterfactual, and prospectively captured, uniformly classified events avoid the selective recall that inflates spontaneous reporting. The scale of the pooled population is also what let a rare outcome be detected at all, since suicidal behaviour is infrequent enough that no single trial could reliably show a difference. The accompanying table sets the two arms side by side on the measures that carry the argument.

TRIAL ARM	PATIENTS POOLED	COMPLETED SUICIDES
Active antiseizure drug	about 28,000	4
Placebo	about 16,000	0

The pooled FDA programme by trial arm: a larger drug-treated population against a placebo comparator, with the completed suicides falling only in the active-drug arms. The event is rare in absolute terms, which is exactly why so large a pooled sample was needed to see it.

The completed suicides are worth holding carefully rather than over-reading. They fell only in the drug-treated groups, which is sobering; but the absolute count is very small, and a handful of events cannot by itself establish a rate. This is why the primary signal was carried by the far more common outcome of suicidal ideation and behaviour, where the numbers were large enough to compare, with the completed suicides serving as corroboration rather than the headline.

26.3 Relative versus absolute risk: reading the magnitude

Roughly doubled in relative terms, but small in absolute terms. The single most misused fact in this topic is the size of the effect, because relative and absolute framings tell very different emotional stories from the same data. On the pooled analysis, suicidal ideation or behaviour occurred **about twice as often on active drug as on placebo, roughly 0.43% versus 0.22%**. Stated as a ratio, a doubling of risk sounds alarming. Stated in absolute terms, the same finding amounts to **about two additional affected patients per 1,000 treated**, which is a modest increment set against the substantial benefit of controlling seizures. Both figures are true and neither is the whole truth; the skill the examiner is testing is the ability to hold them together. Translated into the language of number needed to harm, a great many patients must be treated before one additional suicidality event is attributable to the drug, which is why the warning was framed as a prompt to vigilance rather than a reason to withhold treatment. The point generalises: a relative risk without its baseline is a rhetorical device, and a responsible reading of any safety signal converts the ratio back to absolute terms before acting on it.

RULE **The signal is real, small in absolute terms, and spread across the drug class.** It is an argument for actively monitoring mood and suicidality after starting an antiseizure drug, not an argument for avoiding effective treatment. Convert the relative doubling into its absolute increment before you let it change a prescribing decision.

26.4 Raised across indications, and the confounding problem

The risk held across every indication and was greatest in the epilepsy patients. An obvious objection to any such signal is that antiseizure drugs are prescribed for conditions that themselves carry a raised suicide risk, so that the drugs might merely be markers of a high-risk population. This is the problem of **confounding by indication**, and it is not a hypothetical worry here: these agents are used for migraine, neuropathic pain and bipolar disorder among other indications, each of which is associated with suicidality in its own right. The mood-stabiliser role of the class in bipolar disorder, with its own mechanism and monitoring, is developed in the Mood disorders: pharmacotherapy notes and is not reopened here. The FDA analysis addressed the confounder directly by examining the signal within indication groups, and the finding that gives the warning its weight is that the raised risk was present across all the

indications studied rather than being confined to one high-risk group. Within that spread, the relative increase was highest among the patients treated for epilepsy, which is the group of greatest interest to the psychiatrist working at this interface. That the signal survived the confounding, and was strongest precisely where seizure control is least dispensable, is what makes it a genuine drug effect worth acting on rather than an artefact of who receives these drugs.

TRAP **Attributing an epilepsy patient's whole suicide risk to the drug.** The intrinsic suicide risk of epilepsy and its standardised mortality ratio is a large, separate story taught under its own topic in this chapter, and it is a heavier contributor than the drug signal. The examiner wants the two attributions kept apart: the disease raises suicide risk substantially, and the drug class adds a small further increment on top. Collapsing them, in either direction, loses the mark.

26.5 A class effect across many agents

The warning names a long list, and the class label is deliberate. The breadth of the class effect is easiest to grasp from the agents it covered. In the analysis these were carbamazepine, divalproex sodium and valproate, felbamate, gabapentin, lamotrigine, levetiracetam, oxcarbazepine, pregabalin, tiagabine, topiramate and zonisamide, a set that cuts across chemical classes, mechanisms of action and eras of introduction. That diversity is exactly the point. When a safety signal appears across agents that share little except their antiseizure indication, the parsimonious reading is a class effect rather than a property of any one molecule, and the warning was written to the class for that reason. It would be a misreading to conclude that these agents carry equal psychiatric liability in ordinary practice; individual drugs differ markedly in their propensity to produce low mood, irritability and behavioural change, and those agent-specific psychiatric side-effect profiles, along with the cognitive and behavioural adverse effects of antiseizure drugs more broadly, are the subject of their own topics in this chapter. What the class warning establishes is narrower and firmer: a pooled, class-level increase in emergent suicidality that a prescriber of any of these drugs should have in mind, whichever particular agent is chosen.

IN INDIA

The practical shape of this caution differs here. Much prescribing still rests on older, cheaper agents, follow-up is often hurried, and structured suicidality screening at each visit is far from routine, so the burden of noticing an emergent change often falls on the family. The workable response is to build the warning into psychoeducation from the first prescription: tell the patient and a family member, in plain language and in the patient's own tongue, what an emerging change in mood or thinking looks like and how to raise it quickly, since a WhatsApp message or a call to the treating team is a more realistic safety net than a distant review appointment. Stigma around both epilepsy and suicide makes families slow to volunteer these observations, which is another reason to name them explicitly.

26.6 A warning, not a black box

The regulatory form is the fact most often upgraded from memory. When the joint FDA advisory committees deliberated, they recommended that the labelling of these drugs carry **warnings and precautions** about suicidality, but they stopped short of a **black box warning**, the strongest and most restrictive form of labelled caution. The distinction matters both clinically and for the examination. A black box is reserved for hazards judged severe enough to demand the most prominent possible flag and often to constrain use directly; the committees' judgement that a small absolute excess, set against the clear benefit of seizure control, did not warrant that highest tier is itself the substance of the decision. The warning has also drawn methodological criticism, chiefly that pooling heterogeneous trials with very low event rates and non-uniform ascertainment can generate a fragile signal, and that the class framing may obscure real differences between agents. Those debates do not overturn the core, cautiously stated finding, and the sensible posture is to take the monitoring seriously while not treating the label as a reason to deny patients effective drugs.

TRAP **Calling it a black box warning.** This is the classic upgrade error. The advisory committees voted for warnings and precautions in the package insert and specifically not for a black box. Candidates, reasoning that anything about suicidality must be a black box, misremember the tier and lose an easy mark. Hold the exact regulatory form.

PITFALL

The mirror-image clinical error is over-reaction: withholding or hastily withdrawing an effective antiseizure drug because of the suicidality label, and leaving the patient with uncontrolled seizures. The excess risk is small and absolute, and it is greatest in exactly the epilepsy population who can least afford poor seizure control. The correct response to the warning is to monitor and to educate, not to stop a drug that is working.

26.7 Management

The whole of management follows from the size and shape of the risk. Because the signal is small, class-wide and expressed as emergent change, the response is surveillance and shared vigilance rather than any change in which drug is used for the seizures. The guidelines that frame the work are the epilepsy standards of the International League Against Epilepsy and NICE, which set optimisation of seizure control as the non-negotiable foundation, together with the epilepsy-psychiatry management literature synthesised in Dulcan's Textbook of Child and Adolescent Psychiatry for the handling of the suicidality caution itself. The **LOCUS** of care is overwhelmingly outpatient, since this is a monitoring task layered onto routine epilepsy or psychiatric follow-up; care escalates to emergency or inpatient settings only if active suicidal ideation or intent actually emerges, at which point the management becomes that of any acute suicidality regardless of its trigger. The focus shifts by phase. The acute and short-term task, concentrated around initiation and the early weeks of any new antiseizure drug, is to watch for newly emergent or worsening suicidal thoughts and for the mood and behavioural change that heralds them; the mid-term task is continued periodic review as treatment stabilises; and the

long-term task is to keep the caution alive across a chronic illness without letting it erode adherence to an effective regimen.

The modus spans several modes at once, though none of them is a change of antiseizure agent. The core intervention is clinical monitoring coupled with structured education: clinicians prescribing these drugs should monitor mood and behaviour for emergent suicidality and advise patients and their families to report the warning signs promptly. Making that education concrete is what turns a label into a safeguard, so the specific signs to hand to the household are worth naming.

- New or worsening depressed mood.
- Talking or thinking about self-harm.
- A preoccupation with death or dying.
- Social withdrawal or a marked, unexplained change in behaviour.

The remaining modes are supportive rather than pharmacological in the drug-substitution sense. Psychological input is the safety planning and psychoeducation that make the family an active part of monitoring; the social mode is the family vigilance and accessible contact route that catch an emerging change between appointments; and the separate question of choosing an antidepressant, weighed against its effect on the seizure threshold, is handled in its own dedicated topic in this chapter. The through-line is that the drug keeps doing its job while clinician and family together keep watch for the small, real, detectable risk it carries.

MNEMONIC

FDA: Found pooled, Doubled but small, A warning not a box

The three things to carry are that the signal was **F**ound in a pooled analysis of placebo-controlled trials and holds across the whole drug class, that it represents a **D**oubling in relative terms but only a small absolute excess, and that the regulatory outcome was **A** warning and precaution rather than a black box. The agency's own initials are the memory hook for its own decision.

GOOD TO REMEMBER

Build a one-line suicidality inquiry into the visit at which any antiseizure drug is started and into the early follow-up, the same reflex with which a clinician asks about rash after starting a new agent. Ask about mood and any thoughts of self-harm, document that you asked, and make sure one family member knows the warning signs and how to reach you. It costs a sentence and it is where an emergent change is actually caught.

26.8 Pulling the picture together

One decision, four facts, and a discrimination that must survive. The topic resolves into a compact set of things worth carrying into the examination and the clinic. A drug-safety decision was reached from a large pooled analysis of randomised trials; the risk it identified is roughly a doubling in relative terms and a small increment in absolute terms; it was a class effect across a long list of agents and held across every indication while being greatest in epilepsy; and its

regulatory form was a warning and precaution rather than the black box that memory tends to supply. Above all, the drug-attributable suicidality signal must not be conflated with the far larger, disease-intrinsic suicide risk of epilepsy that is taught under its own topic. Hold that separation, keep the absolute and relative framings side by side, and the management writes itself as monitor, educate, and do not withdraw an effective drug. The at-a-glance strip below fixes the spine.

Class-wide	Small, real	Warning, not a box	Watch, don't withdraw
ALL ANTISEIZURE DRUGS, NOT ONE AGENT	ABSOLUTE EXCESS SUICIDALITY	THE FDA REGULATORY FORM	THE MANAGEMENT STANCE

In 2008 the FDA issued a class warning, not a black box, of increased suicidality across antiseizure drugs, based on a pooled analysis showing suicidal ideation or behaviour about twice as often on drug as on placebo but only a small absolute excess; the risk held across indications and was greatest in epilepsy, and the correct response is to monitor mood and behaviour and educate patients and families, never to reflexively withhold an effective drug or to confuse this drug signal with the disease-intrinsic suicide risk of epilepsy taught elsewhere.

Personality, behaviour and cognition in epilepsy

27 · Temporal lobe epilepsy and interictal personality change

This is the topic where a century of clinical folklore meets a firm modern correction, and the marks lie in holding the tension between them. Interictal personality change refers to the enduring temperament a person with epilepsy carries in the settled periods between seizures - not the fear or déjà vu of the ictus, not the confusion of the postictal state, but the stable, cross-situational way of being that colours ordinary life. For more than a century clinicians believed that chronic epilepsy stamped a characteristic personality on those who had it, and the phrase "epileptic personality" entered the textbooks as though it named a real and unitary thing. The contemporary position is far more careful, and stating it correctly is where the examiner separates the reader who has weighed the evidence from the one who has merely memorised a trait list. The anatomy and semiology of temporal lobe epilepsy, and the named Gastaut-Geschwind cluster, are constructs of their own and taught in their own right elsewhere in this chapter; this topic is the general interictal personality change, the temperamental alteration itself, the instrument built to measure it, and the reasons the old certainty collapsed.

What follows moves from the historical idea and its modern deflation, through the Bear-Fedio inventory, the disputed laterality claim and the specificity problem that unravelled the concept, to viscosity, the social-cognitive substrate, the personality disorders that genuinely cluster with epilepsy, and the markers of risk. Kaplan and Sadock's Comprehensive Textbook of Psychiatry,

Lishman's Organic Psychiatry and Kaufman's Clinical Neurology for Psychiatrists supply the framing between them.

27.1 The idea of an epileptic personality, and why it fails as a general law

A nineteenth-century character that modern evidence dissolved. Early alienists wrote of an epileptic character - adhesive, irritable, egocentric, religiose, slow and circumstantial in thought - and treated it as a near-inevitable consequence of the falling sickness. The modern reading, set out plainly by Kaplan and Sadock, is that there is no single general or specific **epileptic personality** at all: the traits once gathered under that heading are neither universal among people with epilepsy nor uniform in those who show them, and where they do appear they cluster in a subgroup, chiefly those with temporal lobe epilepsy, rather than marking the whole population of patients. That correction does not deny that some patients with long-standing epilepsy change; it denies that the change is one thing, that it is common to all, or that it defines a syndrome. **Interictal personality change**, properly understood, is therefore a description of traits over-represented in a minority, not a diagnosis that any given patient with epilepsy can be expected to develop. The clinical and examinable danger runs in a single direction: the phrase sounds authoritative, it has a long pedigree, and it invites the unwary to expect a personality where there is only a person with a seizure disorder.

RULE **There is no single epileptic personality.** Interictal traits are neither universal in epilepsy nor a unitary syndrome; where they appear they are concentrated in a subgroup, chiefly temporal lobe epilepsy. Every error in this topic comes from promoting a minority description into a general law.

MNEMONIC

Not Universal, Not Specific, Not a Disorder

The three restraints that keep the concept honest: the interictal traits are **not universal** (they appear in a subgroup, not in every patient), **not specific** (the same traits are found in people who do not have epilepsy at all), and **not a disorder** (they are read as markers of brain change, not as a categorical personality disorder). Drop any one restraint and you have reinvented the discredited epileptic personality.

27.2 The Bear-Fedio inventory: an attempt to measure the traits

Turning a clinical impression into a questionnaire. The most influential attempt to give the putative interictal traits an objective footing was the **Bear-Fedio Inventory**, an MMPI-style instrument that Bear and Fedio devised in 1977 to score the specific behaviours long attributed to temporal lobe epilepsy. Where earlier writers had offered adjectives, the inventory offered a scale: it enumerated **18** putative traits - among them circumstantiality, deepened emotionality, a sense of personal destiny and interpersonal viscosity - and asked the patient and an observer to rate each, on the reasoning that a real behavioural syndrome should be visible to the patient and to those around them alike. In their original series of **27** patients with presumed temporal lobe epilepsy, those patients did score higher on the inventory than healthy comparison subjects, and for a time this looked like the objective confirmation the epileptic-personality tradition had

always lacked. The instrument matters for the exam less as a clinical tool - it is not used to diagnose anyone - than as the hinge of the whole debate: it made the traits measurable, and once they were measurable they could be tested against the right control groups, which is precisely how the syndrome came undone.

27.3 Laterality: an emotional right, an ideational left

The claim that the side of the focus shapes the character. One of the inventory's most quoted findings was that the trait profile seemed to differ with the **side of the temporal focus**. Patients with a right temporal focus were reported to show the more overtly emotional traits - a deepened, sometimes sorrowful emotionality and a strong moralising tendency - while those with a left temporal focus showed a more ideational, ruminative profile weighted towards religious and philosophical preoccupation, humourlessness and a sense of personal destiny. The distinction was seductive because it fitted a broader intuition that the right hemisphere is the more affective and the left the more verbal and conceptual, so that a discharging right temporal lobe might exaggerate emotional expression and a discharging left temporal lobe might exaggerate the life of ideas. It is worth carrying as a described association and a favourite of examiners, but it must be carried with its health warning: the laterality claims were later disputed and have not been robustly replicated, and they belong to the same body of findings that the specificity critique went on to undercut. The accompanying figure sets the profiles side by side; the table below holds them for revision.

SIDE OF TEMPORAL FOCUS	TRAIT PROFILE	REPRESENTATIVE FEATURES
Right temporal	overtly emotional	deepened emotionality, a sorrowful affect, hypermoralism
Left temporal	ideational and ruminative	religiosity, philosophical preoccupation, humourlessness, a sense of personal destiny

The disputed Bear-Fedio laterality contrast: emotional traits reported with right temporal foci, ideational traits with left. The association is quoted often but was not robustly replicated.

27.4 Why it is not an epilepsy-specific behavioural syndrome

The control groups that dismantled the concept. The decisive objection to the epileptic personality is one of specificity, and it is the single most examinable point in the topic. When the same rating traits were measured not only in patients with epilepsy but in psychiatric patients without epilepsy and in people with comparable physical disability but no seizures, those comparison groups showed the traits too. A pattern that appears just as readily in the non-epileptic sick and disabled cannot be a signature of epilepsy; it is a marker of chronic illness, brain injury and the psychosocial weight of a disabling condition, distributed across many populations rather than owned by one. Kaufman's is explicit that these are best understood as not an **epilepsy-specific** phenomenon - features of brain damage and its circumstances rather than a distinct, codable personality disorder - and that contemporary video-EEG-based studies have failed to corroborate a discrete epileptic personality. The reasons the concept unravelled are worth listing in order:

- the traits appear in non-epileptic psychiatric patients and in equally disabled controls, so they are not specific to epilepsy;
- they are better read as markers of brain damage and chronic illness than as a categorical personality disorder;
- the rating instruments cannot separate the effect of the epilepsy from the effect of the life lived around it;
- and modern studies with video-EEG confirmation have not reproduced a discrete syndrome.

TRAP **Reifying the trait list into an epilepsy-specific personality disorder.** The candidate who has just learnt the Bear-Fedio traits is tempted to treat them as diagnostic of temporal lobe epilepsy, or as a personality disorder caused by it. They are neither: the same traits appear in controls who do not have epilepsy, and they index brain injury and its burdens rather than a categorical disorder. The examiner is testing whether you know the concept was refuted, not whether you can recite the list.

IN INDIA

Epilepsy already carries a heavy social penalty here - it can end a marriage negotiation, bar a school place and cost a job - so an added label of "epileptic personality" attaches a second, spurious stigma to the first, and it does so on a concept the evidence has discredited. Where access to video-EEG, formal neuropsychology and unhurried longitudinal review is uneven, an interictal trait is especially easily misread as a fixed defect of character. The safer discipline is to describe the behaviour actually observed, to attribute it cautiously, and never to let a refuted label narrow a patient's social options in a setting where the social cost of the word is real.

27.5 Viscosity: the one interictal marker that endures

The trait that survives the critique. If most of the epileptic personality dissolved under scrutiny, one general interictal phenomenon has kept a more robust clinical identity: **viscosity**. Lishman describes it as a tendency to prolonged, repetitive, circumstantial and pedantic social interaction, in which the patient cannot easily bring an exchange to a close. The conversation adheres: the same preoccupations are circled again and again, detail is heaped on detail, the point is approached by a long and winding road, and the patient holds the listener past every natural ending, so that interviews run long and disengagement takes effort. What lifts viscosity above a mere manner is the way it is best understood - not as stubbornness or neediness but as a fundamental abnormality of language and social communication, a disturbance in the pragmatics of talk and turn-taking. Seen this way it overlaps with the communication style described on the autistic spectrum, where the machinery of reciprocal conversation is likewise disturbed, and that framing is more useful clinically than reading the behaviour as a character flaw. Viscosity also appears, under the same name, as one strand of the named Gastaut-Geschwind cluster taught elsewhere in this chapter; here it is the general interictal phenomenon in its own right - circumstantial, over-detailed speech that circles a restricted range of preoccupations, resists closure, and holds the listener with an adhesive, earnest quality that is effortful to leave.

27.6 The substrate: impaired social cognition and theory of mind

Where the interpersonal change may come from. A more mechanistic account of why some patients with temporal lobe epilepsy become interpersonally difficult has grown from the study of social cognition. Kaplan and Sadock report that patients with temporal lobe epilepsy show impaired social cognition and **theory of mind**, with measurable deficits in affective as well as cognitive empathy - the capacity to feel with another and the capacity to model what another is thinking. If a patient reads other minds less accurately and shares their feelings less readily, then the surface behaviour follows: reciprocity falters, conversations misjudge the listener's state, and the interpersonal texture the older literature labelled as personality can be reframed as a consequence of a disturbed social brain. The relevant machinery is the temporolimbic and prefrontal circuitry that supports social inference, sitting close to the very structures that generate temporal lobe seizures, which makes it plausible that chronic disturbance there should show up as altered social behaviour; the anatomy of the limbic system in full belongs to the Functional Neuroanatomy and Neurophysiology notes, and the general neuropsychological and cognitive effects of epilepsy are a separate topic taught in their own right elsewhere in this chapter. What this substrate adds is a kinder and more accurate lens: the behaviour is not chosen, and it is not moral failure.

> **PITFALL**
>
> The common clinical error is to read a temporal lobe epilepsy patient's poor social reciprocity, tangential over-talking and blunted attunement as a difficult or wilful personality, and to respond with irritation rather than accommodation. If impaired social cognition and empathy underlie the behaviour, that response is both unfair and unhelpful: it blames a patient for a deficit they cannot simply switch off. Recognising the social-cognitive basis shifts the clinical stance from exasperation to structured support - clearer communication, more patience with disengagement, and involvement of family who can be taught the same understanding.

27.7 Personality disorder that genuinely clusters with epilepsy

A real comorbidity, distinct from the mythical syndrome. Rejecting the epileptic personality does not mean epilepsy and disordered personality never travel together; they do, and the association is genuine. Kaplan and Sadock report that formal personality disorders are highly prevalent in epilepsy, more so than in the general population, and that the most common of them is **borderline personality disorder**, with dependent and avoidant patterns also over-represented - Lishman puts the excess of these traits at roughly **two-fold** relative to unaffected comparison groups. Keeping this straight turns on a pair of distinctions. First, a categorical personality disorder is a different object from the vague interictal trait description: borderline, dependent and avoidant disorders are recognised diagnoses with their own criteria, and finding them more often in epilepsy is an epidemiological fact, not a resurrection of the epileptic personality. Second, the enduring personality picture must be separated from the interictal dysphoric disorder of epilepsy, the episodic mood syndrome taught in its own right elsewhere in this chapter; a persistent way of relating is not the same as a fluctuating affective state, and conflating them muddles the picture. The clinical value of the comorbidity is that these disorders

carry their own management and their own risks, and they are missed if the clinician is looking only for a mythical epileptic character.

27.8 Who is at risk, and what recognition should change

The recognition changes what you understand more than what you prescribe. Some features of the epilepsy make an interictal behavioural or personality change more likely, and knowing them lets the clinician anticipate rather than discover it. Kaufman's identifies a recognisable risk profile: **early-onset epilepsy**, seizures that secondarily generalise, and bilateral EEG changes, with a superimposed psychosis or mood disorder raising the risk further still. The markers to weigh are:

- an early age at onset of the epilepsy;
- seizures that secondarily generalise rather than staying focal;
- bilateral, rather than strictly unilateral, EEG changes;
- and comorbid psychosis or a mood disorder, which compounds the risk.

There is no specific pharmacotherapy for interictal personality change itself; the point of recognising it is to interpret the patient correctly and to find the treatable disorders that hide behind it. Management therefore rests on the same footing as the rest of epilepsy-psychiatry care, framed by the ILAE and NICE epilepsy guidelines for the seizures and by the epilepsy-specific psychiatric-management literature for the behaviour. The **LOCUS** of care is overwhelmingly the outpatient clinic, a chronic pattern followed over time, reaching inpatient or emergency settings only through a comorbidity - a psychosis, a mood crisis or a safety concern - never through the trait change itself. The **focus** moves with the horizon: the acute task is recognition and differentiation, separating a described interictal trait from a categorical personality disorder, from the interictal dysphoric disorder and from an evolving psychosis; the short-to-mid-term task is to optimise seizure control and to screen for and treat the mood, anxiety and psychotic comorbidities that carry the real morbidity; the long-term task is psychosocial, supporting relationships, work and adjustment around a stable temperament. The **modus** spans every mode without aiming a drug at personality: the pharmacological lever is pointed at seizures and at any comorbid disorder; the psychological and social modes carry most of the work through psychoeducation, family involvement and attention to marital and occupational consequences; and the procedural lever is the optimisation of antiepileptic treatment that serves the epilepsy at large.

GOOD TO REMEMBER

When a patient with long-standing temporal lobe epilepsy strikes you as sticky in conversation, over-serious and hard to read socially, name it to yourself as a described interictal trait, not as an epileptic personality and not as a chosen difficulty. Then do the useful thing the label cannot do: screen deliberately for depression, anxiety and psychosis, because those are prevalent, treatable and far more dangerous than any trait, and they are exactly what a fixation on personality causes you to miss.

No single type	18	Not specific	Borderline
OF EPILEPTIC PERSONALITY	BEAR-FEDIO PUTATIVE TRAITS	TRAITS ALSO SEEN IN CONTROLS	COMMONEST PERSONALITY DISORDER IN EPILEPSY

There is no single epileptic personality: interictal traits are neither universal in epilepsy nor specific to it, are concentrated in a temporal lobe subgroup, and index brain change rather than a categorical disorder; viscosity is the one durable marker, and the clinical priority is to screen for the treatable psychiatric comorbidities the label conceals.

28 · Gastaut-Geschwind syndrome: hyperreligiosity, hypergraphia, viscosity, altered sexuality

This is the eponym examiners love and candidates over-claim. The Gastaut-Geschwind syndrome names a recognisable pattern of behaviour and personality that emerges in a subset of people with long-standing temporal lobe epilepsy - not during a seizure and not in its immediate aftermath, but in the settled periods between fits. It is one of the very few places in neuropsychiatry where a distinctive interictal behavioural signature has been tied, however loosely, to a particular epileptogenic zone, and that is exactly why it earns marks: it forces the candidate to hold two ideas at once, that the cluster is real and clinically recognisable, and that its status as a specific, universal consequence of temporal lobe epilepsy is far from settled. Hold the first half without the second and you sound credulous; hold the second without the first and you sound as if you have never met such a patient.

The syndrome is built from four cardinal features: a tendency to write, an intensified preoccupation with religious and metaphysical themes, a distinctive interpersonal and conversational stickiness, and a change in sexual behaviour that is usually a reduction of drive. Aggressiveness is sometimes added, though aggression, irritability and episodic dyscontrol in epilepsy are treated as their own topic elsewhere in this chapter. Kaplan and Sadock's Comprehensive Textbook of Psychiatry frames the cluster and its laterality, while Lishman's Organic Psychiatry supplies the original formulation and the objective evidence for the writing. What follows takes the syndrome as a whole, then each feature in turn, and closes with what recognising it should and should not change.

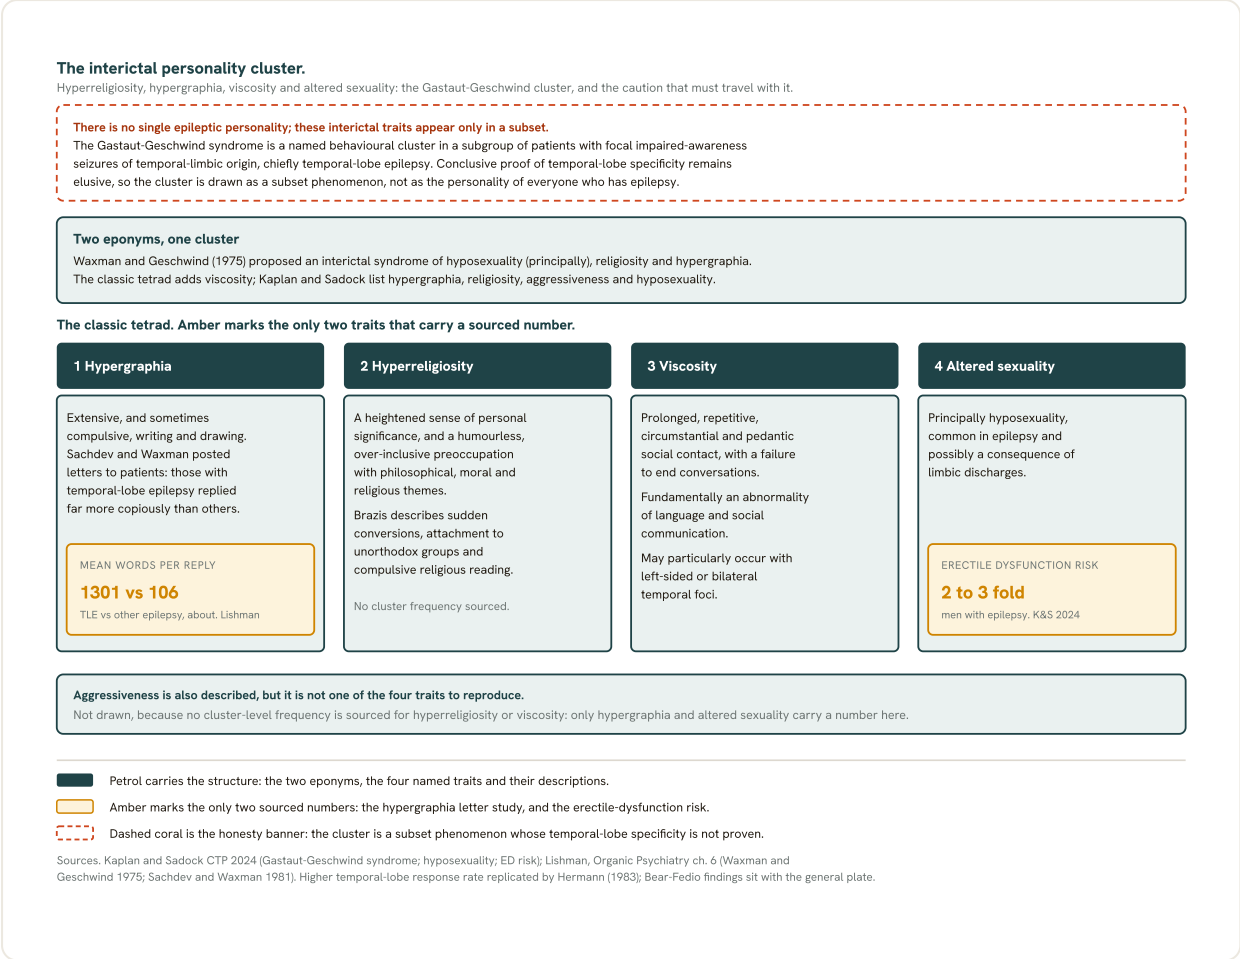


Fig 28.1 The interictal personality cluster. The Gastaut-Geschwind cluster is hyperreligiosity, hypergraphia, viscosity and altered sexuality (principally hyposexuality), seen in a subset of patients with focal impaired-awareness seizures of temporal-limbic origin, not in everyone with epilepsy. Only two of the four carry a sourced number: hypergraphia, shown objectively by Sachdev and Waxman's letter study (mean reply about 1301 words in temporal-lobe epilepsy against 106 in other epilepsy), and altered sexuality, where men with epilepsy carry a two- to threefold risk of erectile dysfunction. Aggressiveness is listed alongside the tetrad but is not one of the four, and conclusive proof of temporal-lobe specificity remains elusive.

28.1 A named interictal cluster, of uncertain specificity

Between the seizures, in some of the patients, from a temporal-limbic focus. The **Gastaut-Geschwind syndrome** is an interictal behavioural cluster arising in a subset of patients whose seizures are focal impaired-awareness seizures of temporal-limbic origin. Three qualifiers in that sentence carry the whole examinable weight. It is interictal: the behaviour belongs to the enduring personality of the patient between attacks, not to the ictus or the postictal state, which is what separates it from the peri-ictal syndromes that occupy much of this chapter. It is a subset phenomenon: the great majority of people with temporal lobe epilepsy never develop the full picture, so the syndrome is neither a diagnostic marker of the epilepsy nor an inevitable outcome of it. And it is tied to a temporal-limbic focus, which is the anatomical claim that makes it interesting and the anatomical claim that remains only partly proven. Kaplan and Sadock are explicit that conclusive proof of specificity to temporal lobe epilepsy remains elusive: the cluster is described most often in this population, but whether it is genuinely specific to it, or a more general accompaniment of chronic epilepsy and its psychosocial burden, has never been settled.

It is therefore best carried as a clinical gestalt with a strong association and a weak proof, not as a fixed lesion-to-behaviour law.

RULE **It is interictal, partial and associated - never ictal, universal or proven.** The Gastaut-Geschwind cluster is a feature of the personality between seizures, seen in only a fraction of temporal lobe epilepsy, and linked to that epilepsy by association rather than by demonstrated specificity. Every examinable error in this topic comes from dropping one of those three restraints.

TRAP **Treating the syndrome as a proven, near-universal stamp of temporal lobe epilepsy.** Candidates who have just learnt the four features are tempted to promote them into a specific, reliable sign - as if a religious, prolific writer must have a temporal focus, or as if temporal lobe epilepsy must produce the cluster. The source itself withholds that certainty: the features occur in a subset and the specificity is unproven, so the inference the examiner is probing runs only weakly, and never as a rule.

28.2 The paradigm: behaviour tied to localisation

An argument about the brain, dressed as a syndrome. The modern formulation belongs to Waxman and Geschwind, in 1975, who described an **interictal behavioural syndrome** whose core they took to be hyposexuality above all, together with deepened religiosity and a tendency to write. What made the proposal influential was less the list than the ambition behind it: they offered the cluster as a paradigm, a worked example of how a circumscribed change in behaviour might be read as the signature of a particular cerebral localisation. When the idea that focal brain activity could sculpt personality was still contentious, an epilepsy that seemed to leave a consistent behavioural fingerprint was powerful evidence for localisationism, and the syndrome has carried that charge ever since. This is also why it attracts as much scepticism as devotion. The very features that make it a neat demonstration of brain-behaviour linkage - that they are qualitative, that they overlap with ordinary human variation, that religiosity and volubility are common enough in people without epilepsy - are exactly what make it hard to prove that the epilepsy, rather than the life lived around it, produced them. State both the claim and its fragility in the same breath, because the examiner's follow-up is almost always whether the evidence holds.

28.3 The four cardinal features as a cluster

Four strands that tend to travel together. Taken together, the syndrome's cardinal features are hypergraphia, hyperreligiosity with a deepened metaphysical interest, viscosity, and altered sexuality that is principally a reduction of drive; Kaplan and Sadock's own enumeration lists hypergraphia, religiosity, aggressiveness and hyposexuality, so aggressiveness recurs as a fifth, less constant member. The clinically important word is cluster. The features are described as tending to co-occur and to shade into one another - the same overinclusive, weighty seriousness runs through the compulsive writing, the moral-religious preoccupation and the conversational adhesiveness - but they do not appear as a fixed tetrad in which all four are always present. A patient may show florid hypergraphia with little religiosity, or a deep religiosity with unremarkable writing. The accompanying figure sets the features around their shared temporal-

limbic substrate; for revision, a four-part memory hook holds them without implying that all four must be present at once.

MNEMONIC

Write, Worship, Weld, Wane

The four cardinal features as four W's: the patient **writes** copiously and often compulsively (hypergraphia), **worships** with intensified religious and metaphysical preoccupation (hyperreligiosity), **welds** to people and to a narrow range of topics in a sticky, hard-to-interrupt way (viscosity), and finds sexual drive **waning** (hyposexuality). The hook is a memory aid only; it asserts nothing about how many of the four any given patient will show.

CARDINAL FEATURE	HOW IT PRESENTS	THE POINT TO HOLD
Hypergraphia	extensive, sometimes compulsive writing and drawing	the one feature with objective, quantified support
Hyperreligiosity	heightened significance and moral-religious preoccupation, at times sudden conversions	hardest to separate from cultural norm
Viscosity	repetitive, circumstantial talk around a restricted range of topics	may favour a left-sided or bilateral temporal focus
Altered sexuality	principally hyposexuality, a lowered drive	has a probable neurophysiological, not only psychosocial, basis

The Gastaut-Geschwind cluster across its four cardinal features. They tend to co-occur and share a weighty, overinclusive quality, but the syndrome is a cluster, not an all-or-none tetrad.

28.4 Hypergraphia: the feature you can measure

The pen that will not stop. Of the four features, **hypergraphia** is the one with a definition sharp enough to be tested and evidence solid enough to be quoted. Lishman describes it as a tendency towards extensive and sometimes compulsive writing and drawing: the patient keeps voluminous diaries, copies out passages, annotates and re-annotates, writes long unsolicited letters, and produces the material with a driven, meticulous quality rather than for any external purpose. The content often carries the same moral or religious weight that colours the rest of the syndrome, which is one of the ways the features are said to cohere. What lifts hypergraphia above anecdote is that it has been shown objectively. In a much-cited study, Sachdev and Waxman, in 1981, sent a simple letter to patients and measured how much they wrote back: those with temporal lobe epilepsy replied far more copiously than patients with other forms of epilepsy, with a mean reply of about **1301** words against about **106** words for the comparison group. The size of that gap is the point - it is not a marginal tendency but an order-of-magnitude difference in output - and the higher response rate was afterwards replicated by Hermann, in 1983. Hypergraphia is therefore the feature to reach for when asked whether any of this can be demonstrated rather than merely asserted.

28.5 Hyperreligiosity and deepened metaphysical interest

Not a delusion, but a change of emphasis. **Hyperreligiosity** here is a heightened, humourless sense of the significance of things, an overinclusive preoccupation with philosophical, moral and religious themes that comes to dominate the patient's inner and outer life. Kaplan and Sadock describe its more florid forms:

- a pervasive sense of heightened significance attaching to ordinary events;
- humourless, overinclusive dwelling on philosophical and moral questions;
- sudden religious conversions, at times more than one over a lifetime;
- attachment to unorthodox or fringe religious groups;
- and compulsive reading of religious texts.

The clinical craft lies in not mistaking this for what it is not. It is a change in the intensity and preoccupying quality of belief, arising in a person whose reality testing is otherwise intact, and it is not, in itself, a psychotic phenomenon; treating a patient's deepened devotion as a delusion to be suppressed both misreads the sign and injures the therapeutic relationship. Equally, the feature must not be inflated into a diagnostic test. Intense religiosity is common in people who have never had a seizure, and the syndrome supplies no threshold above which faith becomes pathological.

IN INDIA

Nowhere is the boundary between the sign and the norm harder to draw than in a society where intense devotion, daily ritual, pilgrimage and scriptural reading are woven into ordinary healthy life. A patient who prays at length, keeps religious diaries and speaks constantly of moral and spiritual matters may be entirely typical of their family and community, not hyperreligious in the syndromic sense. The change that matters is a shift from the person's own baseline - a new, driven, overinclusive preoccupation that those who know the patient recognise as different - and it is read against that baseline and that culture, never against a Western yardstick of ordinary belief. Getting it wrong risks pathologising healthy faith or, in the other direction, missing a genuine change because the content looks culturally unremarkable.

28.6 Viscosity: the interpersonal and conversational stickiness

Hard to start a conversation with, harder to end one. As a feature of the cluster, **viscosity** is the tendency to talk repetitively and circumstantially about a restricted range of topics, so that speech and social contact acquire an adhesive, hard-to-disengage quality. The patient circles the same preoccupations, resists moving the conversation on, and holds the listener in a way that can feel warm and earnest but is effortful to leave; interviews run long and return, again and again, to a narrow set of concerns. Within the Gastaut-Geschwind picture, viscosity is the interpersonal correlate of the same overinclusive, weighty cognitive style that produces the writing and the religiosity. Kaplan and Sadock note a laterality association worth carrying: viscosity may fall particularly to those with a left-sided or bilateral temporal focus, in contrast to features less tied to the side of onset.

Viscosity is also the point at which this topic brushes against its neighbour. The broader interictal personality change of temporal lobe epilepsy - the general alteration in temperament of which conversational and interpersonal stickiness is one described element - is a separate subject taught in its own right elsewhere in this chapter, as is the anatomy and semiology of temporal lobe epilepsy itself. Here the concern is narrower: viscosity as one cardinal strand of the named eponymous cluster, not the general personality change of which it can also be read as a part. Keeping the two apart is what stops this account from swallowing its neighbour.

28.7 Altered sexuality: hyposexuality and its basis

Usually less, not more. The change in sexual behaviour that completes the cluster is, in the great majority, a reduction: **hyposexuality**, a lowered sexual interest and drive, common in epilepsy generally and not confined to the syndrome. What makes it examinable is the argument about mechanism. It would be easy to attribute a flattened libido in a person with chronic epilepsy entirely to the psychosocial weight of the illness, to mood, or to medication, and all of those contribute. But the feature is thought to have a genuine neurophysiological component, possibly a consequence of the same limbic discharges that drive the seizures, so that the temporal-limbic disturbance shaping the rest of the personality also reaches the circuitry of sexual drive. The evidence marshalled for a physical basis is that men with epilepsy carry about a **two- to threefold** increased risk of erectile dysfunction, a raised rate hard to explain by attitude alone and pointing instead to a physiological substrate. The limbic mechanism is named here only as the probable seat of the change; the anatomy of the limbic system in full belongs to the Functional Neuroanatomy and Neurophysiology notes.

> **PITFALL**
>
> The recurring clinical error is to file a patient's lowered libido entirely under antiepileptic side effect or low mood, adjust nothing, and move on. That misses two things: the change may have a neurophysiological basis in the epilepsy itself, and, whatever its cause, it is distressing, treatable and almost never volunteered. Patients rarely raise sexual dysfunction unprompted, so unless it is asked about directly it goes unrecorded and unaddressed. The safer habit is to enquire, to weigh the contributions of drug, mood and disease rather than assuming one, and to treat the problem as real.

28.8 Recognising the cluster and what to do with it

The syndrome changes what you understand more than what you prescribe. There is no specific pharmacotherapy for the Gastaut-Geschwind cluster, and the point of recognising it is largely to interpret a patient correctly and to avoid two opposite mistakes: medicalising a personality into a disease, and dismissing genuine, treatable distress as mere temperament. Management therefore rests on the same foundations as the rest of epilepsy-psychiatry care, framed by the ILAE and NICE epilepsy guidelines for the seizures themselves and by the epilepsy-specific psychiatric-management literature for the behavioural picture, with the accompanying textbook accounts as connective cover.

The **LOCUS** of care is overwhelmingly the outpatient clinic; this is a chronic interictal pattern followed over time, and it reaches inpatient or emergency settings only through a comorbidity - a superimposed psychosis, a mood crisis or a safety concern - rather than through the cluster itself. The **focus** shifts with the horizon. The acute task is recognition and differentiation: placing the hyperreligiosity, hypergraphia, viscosity and hyposexuality against the patient's own baseline and culture, and separating the deepened religiosity of the syndrome from a psychotic illness that would demand a different response. The short-to-mid-term task is to address what is treatable - optimising seizure control, reviewing the contribution of medication and mood to the sexual and behavioural change, and enquiring directly about the hyposexuality the patient will not raise. The long-term task is to support adjustment and relationships around a stable interictal personality, and to watch for the comorbidities that do carry specific treatment. The **modus** spans every available mode without leaning on a drug for the cluster: the pharmacological lever is aimed at seizures and at any comorbid psychiatric disorder, not at personality; the psychological and social modes carry most of the work, through psychoeducation of patient and family, validation of the distress, and attention to the marital, occupational and sexual consequences; and the procedural lever is the same optimisation of antiepileptic treatment that serves the epilepsy at large.

GOOD TO REMEMBER

When a patient with long-standing temporal lobe epilepsy strikes you as intensely religious, endlessly voluble on a few themes, a copious writer and sexually withdrawn, name the pattern to yourself as the interictal cluster it may be, then do the two useful things it points to. Ask directly about sexual function, because hyposexuality is treatable and unspoken, and hold the deepened religiosity as a change in personality rather than a delusion unless independent psychotic features say otherwise. The recognition is worth marks and, more importantly, it stops you from treating a personality as a psychosis.

Interictal	**A subset**	**1301 vs 106**	**2 to 3x**
WHEN THE CLUSTER APPEARS	OF TEMPORAL LOBE EPILEPSY	HYPERGRAPHIA, WORDS PER REPLY	ERECTILE DYSFUNCTION RISK IN MEN

The Gastaut-Geschwind syndrome is an interictal behavioural cluster - hypergraphia, hyperreligiosity, viscosity and hyposexuality - seen in only a subset of temporal-limbic epilepsy and linked to it by association rather than proven specificity; recognise it to interpret the patient and to treat the hyposexuality, never to convert a personality into a psychosis.

29 · Aggression, irritability and episodic dyscontrol in epilepsy

Few topics in this chapter are as misremembered as this one. The lay intuition, and a good deal of medico-legal folklore, holds that epilepsy is a disorder of sudden violence: that a seizure can seize the body into a purposeful attack and then wipe the memory of it clean. The examinable reality is close to the opposite. Aggression is certainly over-represented in people with epilepsy, but it is overwhelmingly a phenomenon of the periods between seizures rather

than of the seizure itself, and on the rare occasions it is directly seizure-related it is crude, brief and undirected. Marks are awarded for holding that distinction steadily, because it is the hinge on which both clinical management and forensic opinion turn: get it wrong and a candidate will either medicalise ordinary interpersonal violence as "epileptic" or, worse, endorse a seizure as the author of an organised crime.

The orderly way to hold the subject is to sort aggressive behaviour by its timing relative to the seizure - the same peri-ictal axis, from preictal through ictal and postictal to interictal, that organises the chapter's other syndromes - and then to set beside it two constructs that recur in the exam and the courtroom: the contested notion of an episodic dyscontrol syndrome, and the irritability that colours the phases around a fit. Kaplan and Sadock's Comprehensive Textbook of Psychiatry supplies most of the phenomenology, Kaufman's Clinical Neurology for Psychiatrists the account of interictal violence, and the forensic-psychiatry literature the vocabulary of dyscontrol and criminal responsibility.

29.1 The central point: aggression is usually interictal, not ictal

The discharge is rarely the culprit. Start from the finding that anchors the whole topic: the great majority of aggressive behaviour seen in people with epilepsy is not the direct expression of **epileptiform activity**, and does not coincide with a seizure at all. It belongs to the interictal state, and it answers to the ordinary determinants of aggression in any population - temperament, social adversity, intoxication, provocation and comorbid mental illness - rather than to the electrical event. This matters because the older literature, and popular belief, ran the other way, imagining a specifically epileptic drive to violence lodged in an abnormal personality. The interictal personality change of temporal lobe epilepsy, and the named Gastaut-Geschwind cluster, are real and are taught in their own right elsewhere in this chapter, but neither is a licence to read a given assault as a seizure equivalent. What aggression there is tends to keep company with two associations in particular: a coexisting psychotic illness, or the impulse-control diagnosis of **intermittent explosive disorder**. Kaplan and Sadock frame both points, and both push the clinician toward the same discipline. Before attributing violence to the epilepsy, look for the interictal cause that is almost always there.

RULE **Aggression in epilepsy is interictal until proven otherwise.** The default attribution for a violent act in a person with epilepsy is not the seizure but the interictal state and its comorbidities, because most such aggression bears no relation to the discharge at all. A seizure-based explanation has to be earned by the timing and the character of the act, never assumed from the diagnosis on the chart.

29.2 Ictal violence: what a seizure can and cannot do

Crude, brief and never planned. When violence genuinely is ictal, generated by the seizure itself, it takes a highly restricted form. **Ictal violence** is not organised, not premeditated and not goal-directed; what actually occurs at the onset of a focal impaired-awareness seizure is confined to **simple violent automatisms** such as spitting, the utterance of profanities and an undirected flailing of the arms, the same fragmentary motor and vocal automatisms that characterise these seizures in general. These are reflex fragments, not actions with an object. Kaplan and Sadock

make the strong version examiners want to hear: coordinated, purposeful violence has rarely if ever been convincingly shown to be a seizure. The mechanistic reason turns a fact to be memorised into one that can be reasoned: the impaired awareness, motor disorganisation and dense amnesia of the ictus are incompatible with the sustained intention, target selection and sequencing a purposeful assault demands. A person cannot at once be lost in the disorganised automatism of a focal impaired-awareness seizure, a seizure type defined in its own topic, and be executing a planned attack on a chosen victim. The clinical and forensic weight of this is large, because it is precisely the plea that "the seizure made me do it" which the examinable facts are built to test and, almost always, to defeat.

MNEMONIC

A seizure spits, it does not plan

Everything ictal violence can do sits in that line. The ictus can throw out a fragment - a reflex, a curse, an undirected swipe - but it cannot supply a target, a sequence or a plan, because the seizure state is disorganised and amnesic. Whenever a violent act was coordinated and purposeful, the discharge did not author it, and the true explanation lies elsewhere on the timeline.

29.3 Postictal violence: resistance, not assault

The struggling patient is trying to get away, not to attack. The second seizure-related window is the postictal one, and the violence that belongs to it has a distinctive quality. **Resistive violence** arises during postictal delirium: as the patient surfaces through a clouded, disoriented aftermath, the attempts by staff or family to help or to physically restrain are misread as an attack, and the patient fights back against the perceived threat. It is nondirected and defensive rather than predatory, and the archetype is the confused patient wrestling free of the hands holding them down. The postictal confusion and delirium in which this happens is itself a peri-ictal syndrome taught in its own right elsewhere in this chapter, and delirium as a general state, with its assessment, subtypes and prevention, belongs to the Stroke, TBI, delirium and focal brain syndromes notes; here the point is only that the disorientation supplies the misinterpretation that drives the struggle. A related and more sustained form of post-seizure aggression can appear as part of a postictal psychosis, the abrupt psychosis after a lucid interval that is covered separately in this chapter, where agitation and disturbed behaviour run a longer subacute course. Recognising postictal violence as fundamentally reactive rather than intentional changes how it should be met, and that is where most wards go wrong.

PITFALL

The commonest way this goes wrong is to read a postictally delirious patient's struggling as deliberate assault and to answer it with force - more hands, tighter holds, a show of physical control. That is exactly the input the disoriented brain misreads, and it escalates the very resistance it was meant to stop. The safer response is to lower stimulation, keep the patient safe with the least restraint that will do, allow the delirium to clear, and sedate rather than overpower if control is genuinely required.

29.4 Interictal violence and its correlates

Verbal and minor, and concentrated in particular patients. Away from the seizure, the aggression that people with epilepsy show is usually low-grade. Kaufman's Clinical Neurology for Psychiatrists describes **interictal violence** as consisting, in the main, of verbal outbursts and minor physical acts rather than serious assault, and, the discriminating observation, as concentrated in patients who are antisocial, psychotic or schizophrenic, or intellectually disabled. The corollary is the part examiners like to probe: violent and non-violent patients with epilepsy do not differ in their seizure type, their EEG or their antiepileptic treatment. What predicts interictal violence, in other words, is the person and the comorbidity, not any feature of the epilepsy, which is the same lesson the ictal and forensic data teach from their own directions. A consistent set of correlates travels with aggression in this population, and, tellingly, none of them is an epileptological variable:

- subnormal intelligence;
- lower socioeconomic status;
- childhood behaviour problems;
- a history of prior head injury;
- and, at the level of the brain, possible **orbitofrontal damage**.

Read as a group these are the familiar risk markers for aggression in general - developmental, social and frontal - reinforcing the message that epilepsy-related aggression largely arrives by the usual routes in a person who also happens to have epilepsy. Where a psychotic illness is the driver, its separation from a primary schizophrenic disorder rests on features developed in the Schizophrenia: aetiology, phenomenology, classification and course notes, and is not reopened here.

RELATION TO THE SEIZURE	CHARACTER OF THE VIOLENCE	ORGANISED AND PURPOSEFUL?	TYPICAL SETTING
Ictal	brief, fragmentary automatisms	No, essentially never	at the onset of a focal impaired-awareness seizure
Postictal	non-directed, resistive struggling	No, reactive to a misperceived threat	during postictal delirium or a postictal psychosis
Interictal	mostly verbal, with minor physical acts	Variable, driven by the person not the fit	in antisocial, psychotic or intellectually disabled patients

Aggression in epilepsy sorted by its timing to the seizure. The single examinable thread is that only the ictal and postictal rows are seizure-generated, both are crude and unplanned, and the organised violence a courtroom cares about fits none of the three.

29.5 Epilepsy, crime and criminal responsibility

Over-represented in prisons, but not more violent. The forensic data close the argument the phenomenology opens. Epilepsy is indeed over-represented behind bars: its prevalence among prison inmates runs at **about two to four times** that of the general population. The naive

inference, that epilepsy therefore breeds violent crime, is exactly what the same data refute. Epileptic prisoners commit no more violent offences than non-epileptic prisoners, and EEG abnormalities do not correlate with violent offending. The over-representation is better explained by the upstream company aggression already keeps - the social adversity, head injury, low intelligence and comorbid disorder catalogued among the correlates - than by any direct criminogenic effect of the seizures. The two halves together are the whole point: a raised prevalence that carries no raised violence. For the witness box this is decisive. An organised, directed violent act cannot be handed to a seizure merely because the accused has epilepsy and an abnormal EEG; the timing must place the act within a seizure or its immediate aftermath, and its character must match the crude, unplanned template that is all a seizure can produce. Absent both, the epilepsy is a comorbidity that travelled to court with the defendant, not an explanation for what they did.

TRAP **Accepting the seizure as the author of an organised crime.** The classic error is to reason from an epilepsy diagnosis and an abnormal EEG straight to ictal responsibility for a purposeful, directed act of violence. It fails on two independent counts: seizure-generated violence is only ever crude and unplanned, and abnormal EEGs do not mark out the violent from the non-violent. Demand the timing and the character before you concede an ictal explanation.

29.6 Episodic dyscontrol syndrome: a contested construct

A forensic label, not an epilepsy diagnosis. The term **episodic dyscontrol syndrome** denotes explosive episodes of sudden, uncontrolled violence for which the person is afterwards apparently amnesic. Its surface resemblance to a seizure - abrupt, stereotyped, out of character, followed by a memory gap - is exactly why it recurs in this chapter and the courtroom, and why it is handled with caution. It is a contested, largely forensic construct rather than a specific epilepsy diagnosis; the forensic-psychiatry literature, Jiloha's Forensic Psychiatry among it, places it among the labels invoked to explain violence and criminal responsibility rather than among the epilepsies proper. The candidate's task is to keep three things apart. The syndrome is not proof of an underlying seizure disorder, however seizure-like the narrative sounds. The apparent amnesia is not the genuine post-ictal amnesia of a fit, and can be hard to separate from the self-serving amnesia claimed after any violent act. And the label overlaps heavily with intermittent explosive disorder, the impulse-control diagnosis for which it often functions as an older synonym. Where the phenomenon is genuine it is usually discussed against a background of temporal lobe epilepsy, whose anatomy and semiology are set out separately in this chapter, but that association is one of clinical interest, not established causation. The sound exam stance is to treat episodic dyscontrol as a description in search of a mechanism, not a settled disease entity.

29.7 Irritability: the affective substrate

The lowered threshold behind the act. Aggression names the behaviour; irritability names the state that lowers the threshold for it, and the two are worth separating because they carry different clinical weight and are managed differently. Irritability, a heightened readiness to react with anger to ordinary provocation, is among the most common and most under-recorded neuropsychiatric symptoms in epilepsy, and it is often more corrosive to family life than the

seizures themselves, because it is continuous rather than episodic and rarely volunteered as a complaint. It is distributed across the same peri-ictal timeline that organises everything else in this topic. It clusters in the hours before a fit, as one strand of the preictal prodrome of dysphoria and irritability that heralds an attack in some patients, a prodrome taught in its own right elsewhere in this chapter. It is a defining component of the interictal dysphoric disorder of epilepsy, the pleomorphic interictal mood syndrome that likewise has its own dedicated topic and that must not be collapsed into an ordinary depressive episode. And it colours the postictal phase in many patients as the storm settles. The connecting message, rather than the syndromic detail that belongs to those owners, is this: irritability is the affective soil from which much epilepsy-related aggression grows, and because patients so rarely offer it, asking about it directly and validating it to the family is often the single most useful thing done in a review that was ostensibly about seizures.

29.8 Assessment and management

What to actually do with the aggressive patient who has epilepsy. Because most of this aggression is not ictal, management is mostly the management of the interictal cause, wrapped around the discipline of first optimising seizure control. The standard of care for the epilepsy itself is set by the ILAE and NICE epilepsy guidelines, while the behavioural disturbance draws on the epilepsy-specific psychiatric-management literature; none of the three licenses a reflex reach for sedation as the whole answer.

The **LOCUS** of care follows the risk and the phase. A florid postictal or psychotic disturbance with a real safety concern may need brief inpatient containment, and an acutely violent patient may first reach an emergency setting; established interictal aggression and irritability, by contrast, are outpatient work, built on careful assessment, treatment of the comorbidity and structured follow-up. The **focus** shifts with time. The acute target is immediate safety and de-escalation; the short-to-mid-term target is the treatable driver, the psychosis, the intermittent explosive disorder, the intoxication or the social crisis, identified by placing the act on the peri-ictal timeline and reconstructing its character; the long-term target is relapse prevention through genuine seizure control and through sustained treatment of the comorbid disorder. The **modus** spans every available mode. The pharmacological lever is aimed at the cause rather than at "aggression" as an abstraction, and where an antipsychotic is used, its selection against the seizure threshold is a decision handled in its own dedicated topic and not reopened here. The psychological and social modes carry much of the interictal work: behavioural analysis of triggers, family education about the reactive nature of postictal struggling, and attention to the adversity and substance use the correlates flag. The procedural lever is the optimisation of antiepileptic treatment itself, prophylactic against the very ictal and postictal windows in which seizure-related violence lives.

GOOD TO REMEMBER

Before you attribute any violent act to a patient's epilepsy, reconstruct two things: when it happened relative to the last seizure, and what it actually looked like. Seizure-related violence is confined to the ictus and its immediate aftermath, and it is always crude and unplanned. If the act was organised, directed and outside those windows, the epilepsy is a bystander, and the assessment should turn to the psychosis, impulse-control disorder, intoxication or social crisis that is doing the real work.

IN INDIA

Two realities sharpen this locally. Stigma still carries the old belief that epilepsy is a disorder of madness and violence, so families, and sometimes courts, over-attribute aggression to the seizures; part of the clinician's job is to correct that on the evidence, since most such aggression is interictal and the epilepsy explains little of it. At the same time the wide treatment gap means poorly controlled seizures, and the clusters that generate postictal disturbance, are common, so optimising antiepileptic treatment is often the single most neglected useful lever. Where a forensic question arises, the same discipline of timing and character applies, and the epilepsy must not become a blanket explanation for a directed crime.

Interictal	**Automatisms**	**Resistive**	**2 to 4x**
WHERE MOST AGGRESSION SITS	THE ONLY TRUE ICTAL VIOLENCE	POSTICTAL, A MISREAD RESTRAINT	EPILEPSY IN PRISON INMATES

Aggression in epilepsy is overwhelmingly interictal and driven by comorbidity, not by the discharge; genuine ictal violence is limited to simple, unplanned automatisms and postictal violence to reactive resistance, so an organised, directed violent act can almost never be laid at the door of a seizure - the single fact that governs both the ward and the witness box.

30 · Cognitive effects of epilepsy and of repeated seizures

The whole of this topic turns on a single act of attribution. A patient with epilepsy who is not thinking clearly hands the clinician three candidate explanations, and the marks here are awarded almost entirely for keeping them apart. The cognition may be impaired by the underlying brain pathology that also produced the epilepsy; by the seizures and the epilepsy syndrome themselves; or by the antiepileptic drugs. The cognitive and behavioural adverse effects of the antiepileptic drugs are taught in their own right elsewhere in this chapter and are deliberately left there. This topic owns the first two limbs, what the disease and the repeated seizures do to cognition, and the reward is for attributing correctly rather than reaching reflexively for the seizures as the author of every deficit.

The second habit to unlearn is the old notion of a specifically epileptic dementia: the belief that epilepsy is a progressively deteriorating disease of the mind, each fit leaving the brain a little poorer than it found it. Lishman's Organic Psychiatry is the anchor text, and the sober picture it assembles is that, for the great majority of people with epilepsy, this is untrue. Cognitive

impairment is real and common, but it is usually mild, usually non-progressive, and frequently present before the first seizure was ever recognised. A short list of named syndromes genuinely does dement, and holding those exceptions in mind is what keeps the general rule honest.

30.1 The myth of epileptic dementia

Uncomplicated epilepsy does not lower intelligence. The cleanest refutation comes from population data rather than clinic samples, which are always skewed toward the severe. Measured across an unselected community population of people with epilepsy, intelligence is normally distributed, and this holds even within the temporal lobe epilepsy subgroup, the very patients in whom folklore expected the greatest ruin. The classical demonstration is the **Isle of Wight study**, from which the conclusion is drawn that a lowered intelligence is not intrinsic to uncomplicated epilepsy. Everything hangs on the qualifier buried in the word "uncomplicated": where intelligence genuinely is low in a person with epilepsy, the low intelligence and the seizures are usually two products of one underlying brain lesion, not a sequence in which the epilepsy eroded a previously intact endowment. The term **epileptic dementia**, when used at all, should be reserved for the handful of syndromes that actually deteriorate, never applied to the ordinary patient whose fits are controlled and whose brain is otherwise sound. A candidate who can state that intelligence in epilepsy is normally distributed, and name the source, has dismantled the commonest misconception in the topic.

30.2 Three sources of impairment, and the task of telling them apart

Disease, seizures, drugs. Because cognition in epilepsy has more than one master, the clinician's first job is aetiological rather than descriptive. Impairment can arise from three broadly separable sources, and a given patient may carry all three at once:

- the antecedent brain pathology, the lesion, malformation or diffuse insult that generated the epilepsy and independently damaged cognition;
- the epilepsy itself, meaning the syndrome, the seizure focus and the burden of repeated seizures acting on the developing or mature brain;
- and the antiepileptic drugs, whose sedative and specific cognitive costs are the subject of their own dedicated topic in this chapter and are not reopened here.

Examiners linger here because the three dissociate on the axes that matter. The drug contribution is dose-related and largely reversible, responding to rationalising the regimen; the disease contribution tends to be fixed or slowly evolving; the antecedent contribution was present from the outset and will not move whatever is done to the seizures. Treating the picture as one thing, and in particular as all the fault of the seizures, is how a reversible drug load or a treatable comorbid depression gets missed while the epilepsy is blamed for damage it never did.

RULE **Cognitive impairment in epilepsy is usually mild, usually non-progressive, and must be attributed before it is treated.** The default is not a dementing process. The clinical task is to apportion the deficit between the antecedent brain pathology, the epilepsy and the drugs, because only that apportionment tells you what, if anything, is reversible and what is fixed.

30.3 The impairment often precedes the first seizure

The most decisive evidence in the topic. If seizures were the engine of decline, cognition should be intact at diagnosis and fall thereafter. The paediatric-onset data show the opposite and are worth carrying verbatim. Children presenting with new-onset epilepsy have academic difficulties far more often than healthy controls, on the order of **about 26 per cent** against roughly **4 per cent**, and these problems are usually identified before the first recognised seizure. That temporal ordering is the whole argument: a deficit that is already present when the epilepsy declares itself cannot have been produced by the seizures that follow. It points instead to a shared antecedent, an early neurodevelopmental disturbance expressing itself both as the seizure disorder and as the learning difficulty, so the two are siblings rather than cause and consequence. This finding, associated with the work of Hermann and colleagues, is the empirical spine of the modern view that epilepsy-related cognitive problems are, in large part, not seizure-driven at all. It also reframes the conversation with families, who assume the fits are damaging the child's schooling and can be reassured that the relationship is more tangled and less bleak than that.

30.4 Impairment is broad, mild and present early in both syndrome types

Not a late complication of refractory disease. The antecedent story is matched at the other end by what is found shortly after onset. Careful neuropsychological study of recently diagnosed patients shows broad but mild cognitive impairments present in both generalised and localisation-related epilepsy close to seizure onset, when patients were assessed on average about ten months after onset. Two features earn marks. First, the impairment is not confined to the chronic, drug-resistant patient of the old textbooks; it is measurable near the beginning of the illness, before years of seizures or polypharmacy could accumulate. Second, it is not the exclusive property of the localisation-related, and specifically temporal, epilepsies; the idiopathic generalised epilepsies, long regarded as cognitively benign, carry their own mild diffuse burden. Taken with the antecedent data, this establishes the shape of the ordinary case: a modest, early, fairly stable impairment reflecting the brain's baseline state rather than a mounting toll. The pattern is a plateau, not a slope, and recognising it stops the clinician over-reading a stable deficit as an unfolding catastrophe.

30.5 The temporal lobe memory signature

Where the epilepsy does leave a disease-intrinsic fingerprint. None of this means epilepsy is cognitively silent, and the clearest disease-intrinsic signature belongs to temporal lobe epilepsy, whose anatomy and semiology are set out in their own topic and not reworked here. Its characteristic neuropsychological profile is one of **selective episodic memory** deficits, a difficulty laying down and retrieving new autobiographical and event memory that stands out against relatively better-preserved function elsewhere. The word selective must not be over-read: memory is to the fore, but impairments can occur across all cognitive domains, so a patient is not guaranteed normal on everything except memory. The emphasis is one of mechanism, the mesial temporal substrate of memory being the tissue most implicated, rather than a promise of isolated amnesia. This deficit is a property of the disease and its substrate, distinct in origin from the diffuse dulling the drugs impose, which is why disentangling the two at the bedside is both

possible and necessary. The interictal personality change that can accompany temporal lobe epilepsy is a separate construct with its own topic and is not a cognitive deficit in this sense.

30.6 Material-specific memory and lateralisation

Which memory fails tracks which side is affected. The temporal memory signature carries a refinement examiners love, because it links neuropsychology to laterality. Memory impairment in temporal lobe epilepsy is **material-specific**: a left-sided focus tends to impair verbal episodic memory, while a right-sided focus tends to impair non-verbal, visuospatial memory. The logic is the standard hemispheric one, verbal material handled by the language-dominant hemisphere and non-verbal by its partner, so the epileptogenic temporal lobe embarrasses the class of memory its side ordinarily serves. Two honest caveats belong with the rule. The verbal side of the dissociation is robust and reproducible; the non-verbal, right-sided side is less consistent and should be stated with more reserve. And the pattern is modulated by language dominance, so the confident left-equals-verbal mapping assumes the typical left-dominant brain and can invert in atypical dominance. Clinically the pattern is more than a curiosity: it helps lateralise the seizure focus in the neuropsychological work-up, a verbal-memory deficit pointing toward the left temporal lobe and a visuospatial one toward the right, complementing the imaging and electrophysiology.

> **PITFALL**
>
> A young person with epilepsy who complains of failing memory and poor concentration is easily written up as showing early evidence of a brain being progressively damaged by seizures. That reading skips the two commonest and most treatable contributors. A sedating or high antiepileptic drug load is dose-related and largely reversible and belongs to its own topic in this chapter; and a comorbid depression, which is frequent in epilepsy and taught separately here, can present as a pseudodementia of poor effort and slowed thinking. Rationalise the drugs and treat the mood before you conclude the epilepsy itself is eroding the brain.

30.7 Repeated seizures and status epilepticus

Does the seizure burden itself accumulate? The title asks specifically about repeated seizures, and this is where the disease-intrinsic account meets its hardest question. The theoretical case that recurrent seizure activity could injure the brain is real: sustained or repeated discharge is metabolically punishing and potentially excitotoxic, and the limbic structures that generate temporal lobe epilepsy are also thought to be the ones most vulnerable to it. The bridging theory here, the kindling model and the neurobiology linking epilepsy to behaviour, is developed as the mechanism step of this chapter and is named rather than reworked, and the full limbic anatomy on which it rests belongs to the Functional Neuroanatomy and Neurophysiology notes. Set against that plausibility is the evidence already assembled: if much of the impairment antecedes the seizures and is measurable near onset, a model in which every fit subtracts a little more cognition is hard to sustain for the ordinary patient. The examinable balance is cautious. Prolonged convulsive **status epilepticus** is a neurological emergency with a genuine capacity to leave lasting injury, on a different footing from a scatter of brief seizures. Transient postictal cognitive impairment, the temporary clouding after a fit, is a distinct peri-ictal phenomenon

owned by the postictal-confusion topic in this chapter, and delirium as a general state belongs to the Stroke, TBI, delirium and focal brain syndromes notes; neither is permanent decline.

TRAP **Reading any cognitive impairment in epilepsy as proof of a progressive, seizure-driven dementia.** The classic error runs straight from "impaired cognition" to "the seizures are damaging the brain", and it fails on the evidence: much of the deficit predates the first seizure, most of it is mild and stable rather than mounting, and a large slice of what does fluctuate is the reversible drug effect or a treatable mood disorder. Progression is the exception, confined to named syndromes, not the rule.

30.8 The syndromes where cognition truly deteriorates

The exceptions that keep the rule honest. A rule stated this firmly needs its exceptions named precisely, because in the examination they are where the discriminating marks sit. Two groups genuinely break the pattern of stable, non-progressive cognition. The first is the epileptic encephalopathies, in which the epileptic activity itself drives a regression: the syndrome of continuous spike-wave activity in slow-wave sleep, whose acquired language variant is **acquired epileptic aphasia**, or **Landau-Kleffner** syndrome, produces specific cognitive deficits including a loss of speech comprehension, and its electrographic hallmark is an almost **continuous spike-and-wave discharge during slow-wave sleep**. The second is the rare progressive myoclonic epilepsies, and here the deterioration is neurodegenerative rather than purely epileptic: **Unverricht-Lundborg disease** and **Lafora body disease** are both associated with a progressive ataxia and dementia, and they stand as the explicit exception to the general rule that epilepsy does not cause progressive decline. The unifying teaching is that when a patient with epilepsy is genuinely and progressively losing cognitive ground, the differential narrows toward this small, named set rather than a vague epileptic deterioration.

MNEMONIC

Two encephalopathies, two dystrophies of the brain: Landau, Unverricht, Lafora

When epilepsy truly dements, the answer is almost always one of these named syndromes rather than the seizures at large. Landau-Kleffner regresses through continuous spike-wave in sleep and takes speech comprehension with it; Unverricht-Lundborg and Lafora body disease march on with ataxia and dementia. Learn the three names as the standing exception to the rule that epilepsy does not progress.

SETTING	COGNITIVE PICTURE	COURSE	ATTRIBUTION
Uncomplicated epilepsy, temporal lobe subgroup included	intelligence normally distributed	static	not intrinsic to the epilepsy
New-onset childhood epilepsy	academic difficulty present before the first seizure	antecedent	shared neurodevelopmental origin
Recent-onset epilepsy, either syndrome type	broad but mild impairment near onset	early, plateau	disease-related, not late
Temporal lobe epilepsy	selective episodic memory deficits, left verbal and right non-verbal	relatively stable	disease-intrinsic substrate
Landau-Kleffner syndrome	loss of speech comprehension	can regress	the epileptic activity itself
Progressive myoclonic epilepsy	progressive ataxia and dementia	progressive	the true exception

Cognition in epilepsy sorted by setting. The single examinable thread is that the top four rows are mild, static or antecedent and are not a dementia, while only the bottom two genuinely deteriorate, and they do so as named syndromes.

30.9 Assessing cognition and what to do about it

Attribution is not an academic exercise; it dictates the next move. Because the reversible contributors are the drugs and the mood, the assessment of the cognitively complaining patient with epilepsy is built to find those first. The standard of care for the epilepsy itself is set by the ILAE and NICE epilepsy guidelines, while the neuropsychiatric layer draws on the epilepsy-specific psychiatric-management literature; none of the three sanctions a fatalistic shrug at "epileptic deterioration".

The **LOCUS** is overwhelmingly the outpatient clinic, where cognition is tracked, the drug regimen reviewed and the family counselled; inpatient and neuropsychology-unit assessment is reserved for the presurgical work-up, a suspected epileptic encephalopathy, or a genuinely progressive picture that raises the named degenerative syndromes. The **focus** shifts with the time frame: the acute task is to exclude and treat the reversible contributors, the sedating drug load and the comorbid depression, both handled in their own topics; the mid-term task is formal characterisation, including material-specific memory testing where lateralisation matters; the long-term task is protecting cognition through good seizure control and the least burdensome effective regimen, with educational and vocational support. The **modus** spans every mode. The pharmacological lever is rationalising rather than adding, trimming the offending agents whose cognitive cost belongs to the antiepileptic-drug topic. The procedural lever is seizure control and, in selected refractory temporal cases, the surgical pathway whose cognitive and psychiatric consequences are covered separately in this chapter. The psychological and social levers carry cognitive rehabilitation, memory strategies, family education that the deficit is usually stable rather than doom-laden, and accommodation at school and work.

GOOD TO REMEMBER

When a patient with epilepsy seems to be getting cognitively worse, walk the reversible causes before the fixed ones. Review the antiepileptic drugs for a sedating or excessive load, screen for and treat depression, and confirm seizures are actually controlled. Only when those are addressed and the decline is still advancing should you escalate toward the small differential of genuinely progressive epileptic syndromes and formal neuropsychological and imaging assessment.

IN INDIA

The folk model that epilepsy is a progressive madness rotting the intellect remains widespread and feeds a stigma touching schooling, employment and marriage. The clinical counter is exactly the evidence of this topic: uncomplicated epilepsy does not lower intelligence, and most cognitive impairment is mild and non-progressive. Two local realities sharpen the work. Access to formal neuropsychological testing is limited outside tertiary centres, so a careful history of function and a few bedside measures often stand in for a full battery. And the treatment gap means many arrive after a long duration of untreated seizures, so optimising control and reviewing an inherited drug regimen for cognitive cost are among the highest-yield things a clinician can do.

Normally distributed	**26% vs 4%**	**Antecedent**	**Rarely progressive**
IQ IN UNCOMPLICATED EPILEPSY	ACADEMIC PROBLEMS BEFORE FIRST SEIZURE	IMPAIRMENT OFTEN PREDATES SEIZURES	OUTSIDE THE NAMED ENCEPHALOPATHIES

Cognition in epilepsy is usually mildly impaired, often before the first seizure, and rarely progressive; the exam task is to attribute the deficit between the antecedent brain disease, the seizures and the reversible drug load, and to reserve genuine decline for the named exceptions - Landau-Kleffner and the progressive myoclonic epilepsies - rather than blaming the seizures for a dementia they do not cause.

31 · Cognitive and behavioural adverse effects of AEDs

Every antiepileptic drug is a bargain. It buys seizure control, and it pays for that control in some coin of cognition or behaviour. This topic owns the cognitive adverse effects of AEDs and the behavioural adverse effects of AEDs, and the attribution logic that separates them from the disease. The examinable skill in this topic is not memorising a toxicology list but learning to read the bargain: when a person with epilepsy loses mental sharpness or becomes irritable, you must separate what the drug is doing from what the disease and its seizures are doing. That act of attribution is the whole point. The cognitive effects of the epilepsy itself and of repeated seizures are taught as their own topic, and the disease-intrinsic aggression, irritability and episodic dyscontrol that arise from the condition rather than from its treatment belong to the behaviour-in-epilepsy topic. Here we own only what the drugs do.

Two adverse-effect streams run in parallel, and the trap is to assume they run together. One is cognitive: psychomotor slowing, inattention, memory difficulty and word-finding trouble. The other is behavioural and psychiatric: dysphoria, irritability, aggression and, less often, frank

psychosis. The agents that head the cognitive list are not the agents that head the behavioural list, and a clinician who carries only one of the two maps will keep swapping a patient off a cognitively heavy drug onto a behaviourally worse one and wonder why nothing improves. Our concern is the cognition-and-behaviour lens and the attribution logic behind it.

31.1 The attribution problem: disease, seizures, or drug

Three possible authors, one presentation. A person with epilepsy who reports poor concentration, mental slowing or a shorter fuse presents a single clinical picture with at least three possible authors, and telling them apart is the examinable core of this topic. The first author is the epilepsy itself, with its underlying lesion and the cumulative effect of repeated seizures, a contribution owned by its own topic and not reopened here. The second is the disease-linked change in temperament and impulse control, the interictal personality change and disease-intrinsic irritability that spring from the condition rather than its treatment, both taught elsewhere in this chapter. The third, and the only one this fragment teaches, is the drug. Because these authors produce overlapping phenotypes, you cannot reason safely from the symptom to its cause; the reasoning has to run the other way, from the history of the treatment to the symptom in front of you.

The tools for that reasoning are temporal and pharmacological rather than diagnostic: a genuine drug effect is time-locked to prescribing changes, dose related and reversible, whereas a disease-driven deficit tracks seizure burden and the underlying pathology. A short structured enquiry sorts most cases:

- did the change begin or worsen when a drug was started or its dose was raised;
- does it rise and fall with the dose;
- does it ease when the drug is lowered or stopped;
- and does its character fit the known profile of the agent in question.

No single answer is decisive, and baseline deficits from the epilepsy muddy every one of them, but taken together they let you assign a probable author and act on it. The discipline matters because the drug is the one author you can often simply remove.

RULE **Before you blame the epilepsy or the personality, account for the drug.** A cognitive or behavioural change in a treated patient is drug-attributable until the timeline says otherwise, because the drug's contribution is dose related, time-locked to prescribing changes and, crucially, reversible. The disease and its seizures are the harder authors to modify; the treatment is the one you can actually test by adjusting it.

31.2 Cognitive adverse effects and how the agents rank

Older is usually heavier, with sharp exceptions. As a broad generalisation the established antiepileptic drugs impose more of a cognitive load than the newer agents, mostly through sedation, psychomotor slowing and dulled attention rather than through any focal amnestic effect. Within that pattern there is a clear worst offender. **Phenobarbital** is the agent most consistently held to carry the worst cognitive effects, a reputation earned through its marked

sedative action and its blunting of arousal, attention and processing speed, and it is this profile that has driven its retreat from first-line use wherever alternatives can be afforded. Ranged behind it, and broadly similar to one another rather than cleanly separable, sit the other long-established agents: phenytoin, carbamazepine, valproate and clonazepam share a comparable middle-ground profile, so that swapping among them for cognitive reasons alone rarely buys much. At the favourable end, **two** of the newer agents stand out. **Lamotrigine** and **gabapentin** carry a comparatively clean cognitive profile against the older drugs, which is one reason they are reached for when preserving cognition is a priority. Newer does not mean cognitively free, though, as the next agent shows.

AGENT(S)	COGNITIVE PROFILE	PRACTICAL NOTE
Phenobarbital	the heaviest cognitive burden of the group	sedation, slowed attention and processing
Phenytoin, carbamazepine, valproate, clonazepam	broadly similar to one another	the older-agent middle ground; little to gain by switching among them for cognition alone
Lamotrigine, gabapentin	comparatively favourable	preferred where protecting cognition matters

The cognitive-burden ranking of the antiepileptic drugs owned here. Topiramate is deliberately left out because its cognitive signature is distinctive rather than a matter of degree, and it is taken up in its own right below.

IN INDIA

This ranking is not an academic ordering here, because the agent with the worst cognitive profile is also among the cheapest and most available. Phenobarbital remains in wide first-line use across resource-limited Indian practice, sustained by its low cost, once-daily dosing and secure supply, long after richer systems have pushed it down the list precisely for its effect on cognition. A dulled, underperforming patient on long-standing phenobarbital is thus a common and correctable sight, the saving on the tablet quietly paid for in schooling, work and daily function. Where an affordable alternative exists, the cognitive burden is a legitimate reason to move.

31.3 Topiramate: a cognitive signature of its own

Not sedation, but words that will not come. Topiramate deserves separate treatment because its cognitive adverse effect is qualitatively distinctive rather than simply severe. **Topiramate** can produce a marked and characteristic disturbance dominated by word-finding difficulty and a loss of verbal fluency, with impairment of attention and memory layered on top. Patients and their families describe it vividly: the sentence that stalls halfway, the ordinary noun that will not surface, the sense of having become slow and inarticulate. This pattern is distinctive enough that an experienced clinician can often suspect the drug from the complaint alone, and it separates topiramate from the diffuse sedative dulling of the barbiturate-type effect. The breadth of the agent's pharmacology, acting on several neuronal systems at once, is the usual explanation for why language and executive function are hit so selectively.

The redeeming feature is that the effect is strongly bound to how the drug is introduced. The word-finding and fluency disturbance is substantially minimised by very slow titration, so that

the same total dose reached gradually is much better tolerated than the same dose reached quickly. That turns a drawback into a prescribing instruction: start low, climb slowly, and a large part of the cognitive cost can be avoided. It also means a patient already impaired on topiramate is not necessarily condemned to lose the drug, since dose reduction and patience may recover a good deal of the lost ground.

PITFALL

The avoidable error is to titrate topiramate quickly, provoke the word-finding and fluency disturbance, and then abandon the drug as intolerable. Much of that impairment is a function of the speed of introduction rather than of the molecule as such, so a rapid climb both manufactures the problem and forfeits an otherwise useful agent. Where the drug is genuinely needed, the answer to emerging cognitive complaints is usually a slower schedule and a lower rate of increase, not immediate withdrawal, and the patient should be warned of the effect in advance so that it is reported early rather than silently endured.

31.4 Behavioural and psychiatric reactions: the agents and the spectrum

A different list from the cognitive one. The behavioural and psychiatric adverse effects of the antiepileptic drugs are governed by a different set of culprits from the cognitive effects, which is exactly why the two maps must be held apart. **Four** agents recur at the head of the behavioural list: levetiracetam, tiagabine, topiramate and vigabatrin are the ones most frequently implicated in adverse psychiatric and behavioural reactions. Two of them, **tiagabine** and **vigabatrin**, earn their place chiefly through behavioural rather than cognitive trouble, while topiramate manages to appear on both lists at once. The essential caveat, and a frequent examination point, is that this list is not a whitelist for the rest: no antiepileptic drug is truly free of psychiatric adverse effects, and the named four are simply where the risk concentrates. The systematic per-drug account of these psychiatric effects across individual agents is developed in the antiepileptic-pharmacology material; the point owned here is the behavioural dimension itself and its place beside the cognitive one.

The severity of these reactions runs along a spectrum, and knowing its shape prevents both complacency and alarm. The great majority of behavioural reactions sit at the milder end, taking the form of **dysphoria**, irritability and anxiety, unpleasant and corrosive to daily life but not dangerous and usually reversible on dose adjustment. The more severe end, including frank psychotic reactions, is genuinely less common, so that a florid drug-induced psychosis is the exception rather than the rule even with the agents most prone to cause trouble. An antiepileptic class effect on mood and suicidality, and the regulatory warning attached to it, is a separate topic in this chapter and must not be folded into the ordinary behavioural spectrum described here.

MNEMONIC

L-T-T-V: the four to watch

The agents most frequently blamed for psychiatric and behavioural reactions line up as **L**evetiracetam, **T**iagabine, **T**opiramate and **V**igabatrin. The letters are worth holding, but so is the caution that comes with them: because no antiepileptic is truly free of behavioural risk, remembering the four must not breed false reassurance about everything left off the list.

TRAP **Assuming the cognitive villains and the behavioural villains are the same drugs.** They are largely not. The agent with the worst cognitive profile is not the agent most associated with behavioural disturbance, and several of the drugs kindest to cognition still carry real behavioural risk. Candidates who collapse the two into a single ranking of good and bad antiepileptics get the individual attributions wrong; the disciplined answer keeps a cognitive map and a behavioural map, overlapping only where a drug genuinely sits on both.

31.5 Levetiracetam: irritability, aggression and mood lability

The behavioural agent every candidate must know. Of the behavioural offenders, levetiracetam is the one that turns up most often in vignettes and in clinics, and its profile is worth learning in detail. **Levetiracetam** is particularly associated with a rise in irritability and overt aggression together with mood lability, a cluster distinctive enough that its emergence in a newly treated patient should point straight at the drug. The disturbance can range from a shorter temper and increased argumentativeness to genuine hostility, and it tends to appear early in the course of treatment. The single most important conceptual point is one of attribution: this is a drug-attributable phenomenon, mechanistically and clinically distinct from the disease-intrinsic aggression, irritability and episodic dyscontrol that arise from the epilepsy itself and are taught as their own topic in this chapter. Conflating the two is a live error, because the two demand opposite responses. Disease-driven aggression is managed by treating the epilepsy and its comorbidities; drug-driven aggression is managed by addressing the drug.

That distinction has direct therapeutic teeth. When irritability and aggression emerge in step with starting or up-titrating levetiracetam, the rational move is to reconsider the drug rather than to escalate it or reach reflexively for an antipsychotic to sedate the behaviour, and the disturbance frequently settles once the offending agent is reduced or replaced. The trap is to read a levetiracetam-induced temper as the patient's underlying nature or an expression of their epilepsy, and so leave the true cause untouched.

GOOD TO REMEMBER

When a patient with epilepsy becomes newly irritable, aggressive or emotionally volatile, the first question is not what is wrong with this person but what changed in the drug chart, and when. Map the behaviour against every recent start and dose increase before ascribing it to temperament or to the epilepsy. If the change tracks a new or up-titrated agent, especially one of the recognised behavioural offenders, the most useful intervention is usually to the prescription itself, not a tranquilliser laid on top of it.

31.6 What amplifies the harm: drug load, titration and folate

The dose and the number of drugs matter as much as the choice. The identity of the agent is only part of the story; the way it is prescribed does much to decide whether its cognitive and behavioural potential is ever realised. Three modifiable amplifiers stand out. The first is total drug load: the more antiepileptic drugs a patient carries and the higher the cumulative dose, the greater the cognitive and behavioural cost, so that a polytherapy regimen assembled over years of incremental additions becomes a burden in its own right, independent of any single agent's profile. The second is the speed of titration, already met with topiramate but general in principle, since a fast climb to a target dose provokes adverse effects the same dose reached gradually may not. The third is a metabolic route worth knowing: antiepileptic drugs can deplete **folate**, and because folate status bears on mood and cognition, that depletion is one more thread tying long-term antiepileptic treatment to low mood and mental sluggishness.

Two boundaries keep this honest. The wider pharmacokinetics of how enzyme-inducing antiepileptics interact with other drugs, folate among the substrates affected, is the province of the antiepileptic-drug interaction topic in this chapter and is not reworked here; the point retained is only that folate depletion is a recognised and relevant consequence, offered as a mechanism to be aware of rather than as a licence to assert a supplementation regimen the owned rows do not carry. What the reader should carry away is that load, speed and folate are levers, and that each is adjustable in a way the underlying epilepsy is not.

31.7 Recognising and reversing it: the management logic

The most reversible harm in the chart. Managing the cognitive and behavioural adverse effects of antiepileptics is unusually rewarding precisely because, unlike most of the neuropsychiatry of epilepsy, the cause is one you put there and can take away. The relevant standards are the ILAE and NICE epilepsy guidelines for the seizure disorder itself, read alongside the epilepsy-specific psychiatric-management literature for the behavioural dimension; none of them endorses tolerating an avoidable drug effect as the price of seizure control.

The **LOCUS** of care is overwhelmingly the outpatient clinic, where recognition, dose review and rationalisation of therapy proceed over successive visits; only a severe behavioural reaction, such as a frank drug-induced psychosis with a safety concern, pushes care briefly toward an inpatient or emergency setting. The focus shifts by phase. The acute task is recognition and attribution, placing the symptom on the drug timeline; the short-to-mid-term task is to act on the offending agent by lowering the dose, slowing an over-rapid titration, or switching from a heavy agent to a cognitively favourable one; the long-term task is to keep the regimen as lean as seizure control allows and to monitor cognition and mood over time; the adverse-event burden of the regimen can be tracked formally with the **Liverpool Adverse Events Profile**, a **19-item** self-report scored **19 to 76** in which a higher total means a heavier burden and which carries no single validated clinical cut-off, so it is read as a continuous measure of change rather than as a threshold. The modus spans every available mode. The pharmacological levers are dose reduction, substitution toward a kinder agent, and attention to folate status where relevant. The single most powerful lever is deprescribing: the **withdrawal of unnecessary antiepileptic**

drugs commonly improves cognition and behaviour, and rationalising an overloaded regimen is among the most gratifying interventions in all of neurorehabilitation. The psychological and social modes wrap around this: forewarning patients of a known effect so that it is reported early, psychoeducation for families who may be mislabelling a drug effect as character, and attention to the driving, work and schooling consequences that cognitive dulling quietly imposes.

Cognition	Behaviour	Favourable	Reversible
WORST ON PHENOBARBITAL	LEVETIRACETAM NOTORIOUS	LAMOTRIGINE, GABAPENTIN	WITHDRAW UNNEEDED DRUGS

The cognitive and behavioural adverse effects of the antiepileptic drugs form two separate maps that do not overlap, with phenobarbital heading the cognitive one and levetiracetam the behavioural, and because the cause is a drug you can adjust this is often the most reversible part of the whole clinical picture: before attributing decline to the epilepsy or the personality, account for the treatment.

32 · Psychiatric outcomes and de novo psychiatric disorders after epilepsy surgery

Epilepsy surgery is offered to stop seizures, but whether the patient counts the operation a success is decided as much by what happens to the mind afterwards as by the seizure diary. This topic owns the psychiatric and cognitive outcome of resective epilepsy surgery, and in particular the **de novo psychiatric disorders** that appear for the first time after operation in a person who did not carry them before. The marks reward a balanced ledger rather than a slogan. On one side, removing an epileptogenic focus, most often by **temporal lobectomy**, tends to reduce psychiatric morbidity across the operated population, most convincingly in those whose seizures are abolished. On the other, a recognisable set of new problems - a short-lived depression, an occasional psychosis, a change in sexual behaviour, and losses in language and memory - can arise after operation, and some of them arise whether or not the seizures were controlled. Holding both halves at once is the examinable skill.

Two framing errors are worth naming at the outset: that a good seizure result guarantees a good psychiatric result, and that any postoperative disorder proves the operation went wrong. Neither survives the data assembled below, drawn from Lishman's Organic Psychiatry and from the neuropsychiatric-epilepsy account in Kaplan and Sadock's Comprehensive Textbook of Psychiatry. Psychiatric and seizure outcome are related but not identical, and the months after operation demand their own psychiatric vigilance.

32.1 The two-sided ledger of surgical outcome

Net benefit, greatest with seizure freedom. The honest headline is favourable. Taken as a whole, temporal lobectomy lowers the psychiatric burden a refractory epilepsy population carries, and the benefit is concentrated in those rendered seizure-free. This makes sense, because much of the depression, anxiety and social disability of chronic epilepsy is downstream of the seizures and of a life organised around them, so abolishing the seizures lifts the load. A candidate must therefore not describe surgery as psychiatrically hazardous on balance; it is, on balance, psychiatrically

helpful, and the strongest predictor of psychiatric improvement is the same variable that defines surgical success, freedom from seizures. The gain is never automatic and never uniform. Superimposed on the population-level improvement is a scatter of new disorders the operated patient did not have before, and these de novo phenomena are the substance of this topic. They do not cancel the net benefit but they qualify it, and they are where the discriminating marks sit. The posture that follows is neither triumphalism nor gloom but structured surveillance: expect most patients to do better, and watch for the minority in whom a new disorder declares itself, often in a predictable window and sometimes despite an excellent seizure result.

RULE **Temporal lobectomy reduces psychiatric morbidity overall and most when seizures are abolished, yet distinct de novo disorders can appear after operation independently of the seizure result.** The net direction is benefit, driven by seizure freedom; but a new depression in particular occurs whether or not seizures recur, so the operation's psychiatric and seizure outcomes must be judged on separate axes.

32.2 De novo postoperative depression, the commonest new disorder

Early, self-limiting, and easy to miss. The single most important de novo syndrome is a **de novo postoperative depression**. About **one-third** of temporal-lobectomy patients develop a depressive disturbance in the early weeks after operation, and its texture is characteristic: low mood coloured by anxiety and by **emotional lability**, the affect brittle and tearful rather than uniformly flat. The timing is itself a teaching point. This is not the slow accretion of a chronic mood disorder but an early, phasic disturbance surfacing while the patient is still adjusting to the operation. It is usually self-limiting, tending to resolve over the following months, with most cases settling by around six months. There is also a lateralising hint examiners like: the disturbance is commoner after right-sided lobectomy, which sits against the naive assumption that dominant-hemisphere surgery must be riskier for mood. Because the depression is early, transient and arrives amid an expected convalescence, it is disarmingly easy to fold into a narrative of understandable adjustment and to leave unmonitored, which is exactly the wrong move given what the next section adds about risk.

MNEMONIC

One in three, in the first few weeks, whether the seizures come back or not

The de novo postoperative depression in one line. Roughly a third of temporal-lobectomy patients, arising early in convalescence, with anxiety and emotional lability, usually self-limiting, a little commoner on the right, and, the sting in the tail, occurring independently of whether the seizures were controlled. Learn it as a disorder to be looked for on a schedule, not stumbled upon.

32.3 Why postoperative depression happens, and who is at risk

A biological event, not merely a disappointment. The mechanism dictates whom to watch. Postoperative depression behaves as a biological consequence of temporal-lobe surgery rather than a psychological reaction to a poor result, and the decisive observation is that it occurs equally in patients with and without seizure recurrence. If it were simply demoralisation at a failed operation, it would cluster in those whose seizures returned; it does not. That dissociation

is the empirical heart of the topic. Risk is not evenly spread: the patient at greatest risk is the one with a prior history of depression, so a careful psychiatric history before operation identifies much of the vulnerable group in advance. Crucially, though, the disorder can arise de novo in a patient with no previous psychiatric history at all, which is what makes universal postoperative monitoring, rather than monitoring only the forewarned, the safe policy. The stakes are raised by mood's dangerous companion: an early, under-recognised depression carries a suicide risk, and early recognition is therefore important. The elevated suicide risk that attends epilepsy more broadly is developed in its own topic and not reworked here, but its shadow falls squarely on this postoperative window. A depression that is expected, time-limited and biologically driven is not a reason to withhold treatment; it is a reason to detect it promptly.

TRAP **Assuming postoperative depression only follows a failed operation, or that seizure freedom protects mood.** Candidates instinctively tie psychiatric outcome to seizure outcome and expect the depressed patients to be those whose fits returned. The examinable fact is the opposite: de novo postoperative depression occurs independently of seizure recurrence, so a seizure-free patient is not thereby a mood-safe patient and must still be screened.

PITFALL

The commonest bedside error is to hear early low mood in a recovering surgical patient, judge it a reasonable adjustment to a big operation, and file it as expected without arranging follow-up. Because the disorder is biological and carries a suicide risk, that charitable reading can be lethal. Treat early postoperative low mood as a disorder to be assessed and tracked, and ask specifically about hopelessness and self-harm rather than accepting a smiling reassurance at a busy clinic.

32.4 De novo and persistent psychosis after temporal lobectomy

Uncommon, but it can surface late. Psychosis is the more feared, though less frequent, de novo outcome. In temporal-lobectomy series a psychosis occurs in about **7 to 8 per cent** of patients, and the feature that earns marks is its timing: it can appear or persist even long after the seizures have been arrested. This defeats the tidy expectation that removing the focus should retire the psychiatric risk with the seizures. A **de novo psychosis** emerging months or years into a seizure-free course is a real entity and a reason that psychiatric follow-up after epilepsy surgery cannot be a brief affair closed once the seizures settle. Three boundaries must be kept clear so the reader does not conflate distinct constructs:

- it is not the schizophrenia-like psychosis of epilepsy, the chronic interictal psychosis that has its own topic;
- it is not forced normalisation or alternative psychosis, where disturbance emerges as seizures and the electroencephalogram paradoxically improve, again its own topic;
- and at the bedside it must be separated from a primary psychotic illness, whose phenomenology and course belong to the Schizophrenia: aetiology, phenomenology, classification and course notes.

What this topic owns is the observation that resecting a temporal focus does not immunise against psychosis and may, in a small minority, be followed by one - a reason for sustained rather than perfunctory surveillance.

32.5 Sexual outcome after temporal lobectomy

An outcome that usually improves, if the seizures are controlled. Sexuality is the outcome most often mishandled in the exam because it runs against expectation. The baseline is that most patients with temporal lobe epilepsy are **hyposexual** before operation, that is, they have reduced sexual interest and drive as part of the interictal picture; the altered sexuality of the interictal personality, and its place in the Gastaut-Geschwind cluster, is taught in its own topic and is only the pre-surgical starting point here. What surgery does is, on balance, favourable. After temporal lobectomy, **nearly one-third** of patients show an increase in libido, an unshackling of a previously suppressed drive, and this gain is contingent: it is seen provided the seizures are controlled. That conditionality is the whole teaching. The libido increase is not a random side effect but part of the general improvement that follows seizure freedom, so it belongs with the net-benefit story rather than with the de novo hazards. The practical implication is to counsel patient and partner that sexual interest may rise after successful surgery, a change usually welcome but occasionally destabilising for a relationship organised around a hyposexual partner.

32.6 Cognitive outcome: memory and word-finding

The preoperative brain writes the postoperative report. Cognition is where the surgical trade-off is most explicit, because the tissue removed to stop the seizures is the tissue that serves memory and, on the dominant side, language. The governing principle is that memory outcome is determined primarily by the preoperative state. A significant new memory deficit follows operation in about **5 per cent** of patients in controlled study, and the counter-intuitive refinement examiners prize is who those patients are: the deficit is more likely in those with high preoperative memory function. The logic is that a well-functioning temporal lobe still doing real memory work has more to lose when resected, whereas a diseased, already failing lobe was contributing little, so its removal costs less and can even release the healthier contralateral side. Consistent with that, improvement is possible in those with low preoperative function. The material-specific structure of memory in temporal lobe epilepsy, the verbal-versus-nonverbal dissociation, is a property of the disease taught in its own topic and not reopened here; what this topic owns is the surgical prediction rule that the preoperative baseline, not the operation alone, forecasts the memory result.

Language carries its own postoperative signature. Minor **word-finding difficulty** is relatively common after temporal lobectomy, concentrated after dominant-hemisphere operations, which is unsurprising given that the resection encroaches on language-serving cortex. This is usually a mild anomia rather than a gross aphasia, but frequent enough that it should be named in the consent conversation for a dominant-side procedure rather than sprung on the patient afterward. Serious neurological complications, by contrast, are rare: a major event such as a stroke or a large visual-field defect occurs in under 2 per cent of operations, so the honest framing is of a common

minor cognitive cost, a memory trade-off keyed to the patient's own baseline, and a small but real chance of a major complication.

DE NOVO OR ALTERED OUTCOME	FREQUENCY OR DIRECTION	TIMING AND DEPENDENCE ON SEIZURE RESULT	KEY PREDICTOR
Postoperative depression	about one-third	early weeks, self-limiting; independent of seizure recurrence	prior depression; can be de novo
De novo or persistent psychosis	about 7 to 8 per cent	may appear long after seizures arrested	distinct from schizophrenia-like psychosis
Sexual outcome	nearly one-third increase libido	postoperative, only if seizures controlled	preoperative hyposexuality lifts
New memory deficit	about 5 per cent	determined by preoperative state	high preoperative function at risk
Word-finding difficulty	relatively common	after operation, persistent but minor	dominant-hemisphere resection
Major neurological complication	under 2 per cent	peri-operative	stroke or large field defect

Outcomes of temporal lobectomy sorted by type. The examinable thread is the middle column: depression and psychosis are independent of the seizure result, whereas the libido gain is contingent on seizure control, and the cognitive costs are keyed to the preoperative brain.

32.7 The burden of normality and psychosocial adjustment

Getting better can be its own crisis. Beyond the discrete disorders lies a subtler outcome that connects them. A patient whose seizures vanish must rebuild an identity and a set of relationships that had been organised for years around being a person with epilepsy, and this readjustment can be genuinely destabilising even when, by every clinical measure, the operation succeeded. The phenomenon is sometimes called the **burden of normality**: the paradoxical difficulty of adapting to health, in which a cured patient struggles with new expectations to work, to drive, to take on adult responsibilities and to occupy a well role that family dynamics may resist ceding. It is not a distinct diagnosis and carries no number to memorise, but it is the connective tissue that makes the de novo depression, and the sexual and vocational changes, intelligible as parts of one adjustment rather than unrelated complications. Recognising it shifts the clinician's stance from surprise to anticipation: the postoperative task is not only to abolish seizures but to accompany the patient through the reorganisation of a life, and to normalise the fact that improvement itself can feel like loss.

32.8 Presurgical assessment and postoperative management

The whole topic converges on a single service design: screen before, and watch after. Because the most dangerous de novo outcome is an early depression that is biological, common and independent of the seizure result, the management of psychiatric outcome is built around prediction and surveillance rather than reaction. Three guideline sources anchor the approach without prescribing a numerical protocol: the epilepsy standards of the International League

Against Epilepsy and of the National Institute for Health and Care Excellence set the frame for comprehensive presurgical evaluation and follow-up, and the epilepsy-specific psychiatric-management literature supplies the neuropsychiatric layer neither epilepsy guideline fully covers. Within that frame the LOCUS -> FOCUS -> MODUS structure organises the work.

The **LOCUS** of care is overwhelmingly the specialist epilepsy-surgery centre and its outpatient follow-up clinic, with two flanking settings: before operation, the multidisciplinary presurgical assessment unit, where a formal psychiatric history identifies the prior-depression group at greatest risk and documents any active disorder; and, for a severe postoperative depression with suicidal intent or a frank psychosis, escalation to inpatient or emergency care. The **focus** tracks the phase. The presurgical task is prediction and counselling, mapping vulnerability and preparing patient and family for the burden of normality and the realistic cognitive trade-offs. The acute and short-term task, across the vulnerable early weeks, is active screening for and treatment of the de novo postoperative depression and vigilance for suicidal thinking. The mid-term task is characterising any new cognitive deficit, particularly memory and word-finding, and supporting adjustment. The long-term task is a follow-up long enough to catch a psychosis that may surface well after the seizures have gone. The **modus** spans every mode: the pharmacological mode treats an emergent depression or psychosis on its merits, with the attention to seizure threshold developed in the antidepressant and antipsychotic topics of this chapter; the psychological mode carries pre-operative preparation, monitoring and support through the re-adjustment; and the procedural and social modes are the surgery itself, consented with its psychiatric and cognitive outcomes made explicit, plus the vocational, driving and family support that make the seizure result liveable.

GOOD TO REMEMBER

Two habits change outcomes. First, take a proper psychiatric history before operation, because a documented past depression flags the highest-risk patient and lets you plan follow-up rather than react to a crisis. Second, schedule active mood screening in the first postoperative weeks for every patient, not only the forewarned, and specifically ask about suicidal thinking, because the de novo depression is common, biological and blind to how well the seizures were controlled.

IN INDIA

Comprehensive epilepsy surgery is concentrated in a handful of tertiary centres, and full presurgical evaluation, with prolonged video-electroencephalography, imaging and neuropsychology, reaches only a small fraction of patients with drug-resistant epilepsy. The consequence is twofold. Most eligible patients never reach operation, so the treatment gap sits upstream of the outcomes discussed here; and among those operated, structured psychiatric follow-up can be patchy once the patient returns to a distant district. A realistic plan makes the presurgical psychiatric assessment non-negotiable while the patient is still at the centre, and hands the family and the referring clinician a clear written schedule for postoperative mood monitoring, since the early depressive window will otherwise pass unobserved far from the surgical team.

~one in three	7 to 8%	~one in three	~5%
DE NOVO POSTOPERATIVE DEPRESSION	PSYCHOSIS AFTER TEMPORAL LOBECTOMY	LIBIDO RISES, IF SEIZURES CONTROLLED	SIGNIFICANT NEW MEMORY DEFICIT

Temporal lobectomy reduces psychiatric morbidity overall and most when seizures are abolished, yet a de novo postoperative depression in about a third of patients arises early and independently of the seizure result, an occasional psychosis can surface long after the fits have stopped, libido tends to rise only when seizures are controlled, and cognitive outcome is written mainly by the preoperative brain - so screen before, and watch after.

Psychogenic non-epileptic seizures

33 · PNES: definition, terminology and classification within functional neurological disorder

A name is a clinical act. Few conditions in neuropsychiatry have been renamed as often, or as anxiously, as the paroxysmal episodes that look like epileptic seizures yet arise without the electrical signature of epilepsy. The words a clinician writes in the notes shape how the patient understands the illness, whether the referral is taken seriously, and whether treatment ever engages the psychological mechanism at the root of the problem. For the examination this is the deceptively simple part of the syndrome, the ground on which everything else is built: what these events are, what to call them, and where they sit in the two diagnostic manuals. It rewards precision because it is precisely where loose thinking shows. A candidate who cannot define the entity cleanly, who reaches for a discredited label, or who cannot say whether the disorder is coded as conversion or as dissociation, has signalled that the rest of the answer will be shaky too.

The scope here is deliberately narrow. This topic covers the definition, the terminology and the classification of the syndrome within the functional neurological disorders. The bedside features that separate these attacks from epileptic seizures, the confirmatory investigations, the psychiatric comorbidity and predisposing factors, the delicate task of delivering the diagnosis with its stepped management, and the difficult patient who has both dissociative and epileptic attacks at once are each developed in their own right elsewhere in this chapter. The broad category of functional neurological and conversion disorder, of which this is one presentation, is taught in the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes; here the concern is only where the seizure form of that category belongs and what to call it.

33.1 What the syndrome is: the working definition

A seizure-like event with no epileptiform correlate, presumed to be psychological.

Psychogenic non-epileptic seizures are sudden, involuntary episodes of altered movement, sensation or awareness that outwardly resemble epileptic seizures but are not produced by the abnormal, hypersynchronous cortical discharge that defines epilepsy, and that are presumed to arise from a psychological process. The definition has three moving parts, and each carries clinical weight. The first is the resemblance: the events mimic seizures closely enough to be

mistaken for them, which is exactly why they generate referrals to epilepsy and psychiatry services and why the bedside distinction becomes such a demanding skill. The second is the critical negative: there is no ictal epileptiform activity accompanying the attack, so the event is not an electrographic seizure however convulsive it looks. The third is the positive attribution: the mechanism is presumed psychological, understood in contemporary models as a dissociative one, an involuntary disruption of the normal integration of movement, sensation and awareness rather than a performance the patient chooses to stage. The word presumed is doing honest work here, because the mechanism is inferred rather than directly observed, and the word involuntary is the pivot of the whole topic. What follows from these three parts is that the diagnosis is a positive identification of a functional disorder, not merely the empty space that remains once epilepsy has been excluded.

RULE **This is a positive diagnosis, not a diagnosis of exclusion.** The syndrome is a definable functional disorder in its own right, with its own presumed mechanism and its own place in the classifications, not simply the residue of a negative epilepsy workup. Framing it as an absence, as merely not epilepsy, is the conceptual error that then generates a cascade of downstream problems, from a demoralised patient to a stalled referral.

33.2 A contested vocabulary: from hysteria to neutral terms

The old labels were abandoned because they insulted the patient. No entity in this chapter has attracted more terminological baggage, and the reason is historical. The syndrome descends conceptually from the nineteenth-century notion of hysteria, and its oldest clinical labels carry that inheritance. The terms **pseudoseizures** and hysterical seizures are now regarded as pejorative and have been largely abandoned, and the objection to them is not merely one of fashion. The prefix pseudo implies that the events are false, feigned or somehow not real, which both misrepresents the involuntary nature of the disorder and communicates contempt to a patient who is genuinely suffering. Hysterical carries an even older weight of dismissiveness and a gendered connotation that alienates patients and colours the clinical encounter before it begins. Against this backdrop the field has converged on neutral alternatives. The syndrome is now treated as synonymous with the terms **dissociative seizures** and functional seizures, and the two roots emphasise different things. Dissociative foregrounds the presumed mechanism, aligning the label with the dissociative model of an involuntary interruption of integrated function. The other root emphasises disturbed function of a structurally intact nervous system, sidestepping any statement about mechanism at all. Both were chosen precisely because they describe the disorder without demeaning the person who has it, and both mark a deliberate move away from a vocabulary that pathologised the patient rather than the pathology.

33.3 Why the negative label is unsatisfactory

Defining a thing by what it is not tells the patient nothing. Even the ostensibly neutral term at the head of this topic has a weakness that is worth naming, because it is a recurring source of trouble. The labels built on non-epileptic define the condition entirely by what it is not, and a patient who is told only that the attacks are non-epileptic may reasonably conclude that the doctors have ruled out epilepsy but still do not know what is actually wrong. That reading

corrodes the therapeutic relationship at the very moment it needs to form. For this reason the term **functional seizures** is increasingly preferred as a neutral description that patients find acceptable and that does not prejudge whether the process is conscious or unconscious. The problems with a purely negative framing can be set out plainly:

- it defines the disorder by negation, so it names an absence rather than a presence;
- it reads to the patient as a non-diagnosis, as if the assessment has ended in a shrug;
- it says nothing about the mechanism, and therefore offers no foothold for explanation or for treatment;
- it can be heard as an accusation that the attack was not a real seizure, reviving the very insult that the neutral vocabulary was meant to retire.

The practical resolution, which the wider neuropsychiatric literature has increasingly endorsed, is to lead with a positive functional formulation rather than an exclusionary one: to say what the disorder is and how it works, in a vocabulary that neither blames nor mystifies the patient.

33.4 Events and spells: softening the descriptor further

Sometimes the safest early move is to drop the word seizure altogether. Some clinicians go a step beyond the neutral synonyms and avoid the word seizure entirely in the earliest phase of assessment. Describing the phenomena as **events or spells** is a way of reducing the confusion that arises when a patient hears the word seizure and immediately assumes epilepsy, with all the associations of brain disease, driving restrictions and lifelong medication that the word carries. This descriptive language is most useful while the diagnosis is still provisional, before video-electroencephalography or the settled clinical picture has fixed the label, because it lets the clinician talk about the attacks factually without committing prematurely to a mechanism or importing the wrong frame. It is a small linguistic device, but it reflects the same principle that runs through the whole of this topic: the words come first, and they either open a door to understanding or quietly close one.

> **GOOD TO REMEMBER**
>
> In letters, discharge summaries and the case notes, prefer functional seizures or dissociative seizures as the working term, and reserve the non-epileptic phrasing for the specific context of contrasting these attacks with epilepsy. While the diagnosis is still being settled, events or spells is a safe, honest descriptor. Retire pseudoseizures and hysterical seizures from your vocabulary entirely, including from the notes, because patients read what is written about them and a careless label can undo months of careful work.

33.5 Locating the syndrome in the classifications

Conversion in one manual, dissociation in the other. The two major diagnostic systems agree that this is a functional neurological disorder but reach it by different routes, and knowing both, along with the codes, is a reliable examination point. In DSM-5 the syndrome is classified as a subtype of **functional neurological symptom disorder**, the diagnosis also known as conversion disorder, which sits among the somatic symptom and related disorders. DSM-5-TR keeps that

placement and marks the seizure presentation with a symptom-type specifier, coded **F44.5** and designated **with attacks or seizures**. ICD-11 takes the alternative conceptual route, classifying the disorder among the dissociative rather than the somatic-symptom disorders: it appears as a **dissociative neurological symptom disorder**, with non-epileptic seizures, under the code **6B60.4**. The divergence is not trivial pedantry. It reflects a genuine and unresolved conceptual tension over whether the syndrome is best understood as a conversion phenomenon or a dissociative one, a tension that the two neutral synonyms in common use quietly mirror. Services and examinations still working with ICD-10 will find the entity among the dissociative and conversion disorders of that older system. The general construct of conversion and functional neurological disorder, its mechanism and its wider family, belongs to the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes; the point to hold here is simply where the seizure subtype is filed in each manual and under what code.

F44.5	6B60.4	late teens to early 20s
DSM-5-TR CONVERSION DISORDER, WITH ATTACKS OR SEIZURES	ICD-11 DISSOCIATIVE NEUROLOGICAL SYMPTOM DISORDER, WITH NON-EPILEPTIC SEIZURES	TYPICAL AGE OF ONSET

33.6 Not conscious simulation: the boundary with feigning

The single boundary most often blurred is the one that matters most. The defining feature of the syndrome, and the source of its commonest misunderstanding, is that the attacks are not deliberately produced. This is what separates it sharply from the conditions of conscious feigning, and the separation runs along an axis of awareness and motivation. In **factitious disorder** the illness is wilfully and consciously simulated, the patient deliberately producing or feigning symptoms without any external reward, the motivation lying in the psychological need to occupy the role of the sick person. In **malingering** the feigning is equally deliberate but is driven by a tangible external gain, such as financial compensation, avoidance of work or duty, or the evasion of a legal consequence. The dissociative attack sits apart from both, because the events are not consciously generated at all: the patient is not choosing to produce them, and there is no external incentive to be uncovered. It is useful to think of a spectrum of conscious awareness, with the deliberate deceptions at one end and the involuntary dissociative attack at the other, and with self-deception occupying the uneasy middle ground where a patient is neither fully feigning nor fully aware. The accompanying table sets the three conditions side by side on the axes that separate them.

AXIS	DISSOCIATIVE (FUNCTIONAL) SEIZURES	FACTITIOUS DISORDER	MALINGERING
Production of the episode	Involuntary, not deliberately generated	Consciously and wilfully simulated	Consciously and deliberately feigned
Conscious intent to deceive	Absent	Present	Present
Motivation	Psychological, outside awareness	To occupy the patient role, no external reward	An external, tangible gain

The conscious-awareness axis distinguishing dissociative seizures from the two disorders of feigning; the essential point is that the dissociative attack alone is not consciously produced.

TRAP **Equating a dissociative attack with malingering or putting it on.** Candidates, and colleagues on a busy ward, slide from not epileptic to not real to putting it on, and so treat the patient as a feigner. That is a category error: the attacks are involuntary, and conflating the disorder with factitious illness or malingering both fails the examination and, in the clinic, insults and abandons a genuinely unwell person.

33.7 Who develops it, and where this topic hands over

A young-adult onset, though the disorder respects no age. The syndrome most characteristically begins in the **late teens or early twenties**, the years in which the developmental, interpersonal and psychological stresses that predispose to it tend to converge. That said, onset is not confined to this window, and the attacks can begin across the lifespan, in older adults and in children as well as in the young-adult peak. This demographic sketch is where the present topic deliberately stops. The psychiatric comorbidities that so often accompany the disorder and the predisposing and precipitating factors that make a particular person vulnerable are the substance of a separate topic, as are the mechanics of confirming the diagnosis, the way it should be communicated and managed, and the prognosis, including the outlook for patients who carry both dissociative and epileptic seizures together. Naming those boundaries matters, because a great deal of examinable and clinically important material about this syndrome lives just outside the definition, and the disciplined answer teaches the definition and terminology fully while pointing, rather than straying, to the rest.

33.8 Why the terminology carries the whole case

Get the name right and the treatment has somewhere to stand. It can look, from a distance, as though the argument over what to call this disorder is a semantic quarrel of interest only to nosologists. It is not. The name is the interface between an invisible mechanism and a frightened patient, and it determines whether the explanation that follows will be believed. A label that implies fakery invites denial and disengagement; a label built on an absence invites the conclusion that nobody knows what is wrong; a neutral, mechanistic term that the patient can accept opens the door to a positive formulation and to the psychological work that actually helps. The whole of this topic, the clean definition, the retirement of the pejorative vocabulary, the preference for a functional over a negative term, the correct placement in the classifications, and the firm line against the language of feigning, converges on a single practical goal: to hand the patient an account of the disorder that is both true and bearable. The detail of how that account is delivered and how treatment is then structured belongs to a separate topic; the foundation it rests on is laid here, in the words.

IN INDIA

These attacks are among the commonest functional presentations seen in Indian general and neurology clinics, where they are widely known by lay terms such as fits and are frequently understood by families through idioms of possession, spirit affliction or the older language of hysteria. Many patients reach a psychiatrist only after a long circuit through faith healers, general practitioners and neurologists, often already committed to antiepileptic drugs and to an epilepsy identity. The stigma attached both to epilepsy and to any psychiatric attribution is heavy, and it bears directly on marriage prospects, employment and standing in the family, which makes the choice of a neutral, non-pejorative, family-acceptable term not a courtesy but a clinical necessity. Language that blames or belittles the patient, in a setting where a psychiatric label is already feared, will simply end the consultation.

PITFALL

Recording the diagnosis as non-epileptic or, worse, as pseudoseizures, and stopping there. An exclusionary label committed to the notes travels with the patient through every future referral, propagating an absence in place of a diagnosis and quietly reasserting the old insult. The remedy is to commit a positive functional formulation to the record, in neutral terms, so that the next clinician inherits a diagnosis rather than a negation.

The syndrome is an involuntary, seizure-like event with no epileptiform electroencephalographic correlate and a presumed psychological mechanism; call it functional or dissociative seizures rather than the discredited pseudoseizures, classify it as a conversion or dissociative neurological symptom disorder, and never mistake it for conscious feigning.

34 · Clinical features differentiating PNES from epileptic seizures

Getting this distinction wrong commits a patient to the wrong illness for years. A large proportion of people with psychogenic non-epileptic seizures carry an epilepsy label and swallow antiepileptic drugs for a long time before anyone questions it, and the reverse error, dismissing a genuine seizure as merely functional, is every bit as damaging. The reference standard that settles the matter is prolonged video-electroencephalographic monitoring, which captures an attack together with its brain electrical correlate, but that resource sits in tertiary units and is rarely at the bedside where the decision first has to be made. So the examinable skill, and the daily clinical one, is reading the attack itself: knowing which features of onset, movement, eyes, awareness and recovery push the odds toward a dissociative attack and which toward an epileptic one, and holding firmly to the discipline that not one of them decides the case on its own.

What follows is only the clinical differentiation. What a psychogenic non-epileptic seizure actually is, the terminology it travels under and its place among the functional neurological disorders belong to the companion topic and, more broadly, to the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes. The investigations that confirm it, video-EEG capture and serum prolactin among them, are treated in their own right, as are the psychiatric

comorbidities and predisposing factors that make a person vulnerable, the delicate business of communicating the diagnosis and its stepped management, and the genuinely difficult patient who has both dissociative and epileptic seizures at once. Here we stay with semiology: the phenotype of the event as a witness sees it unfold.

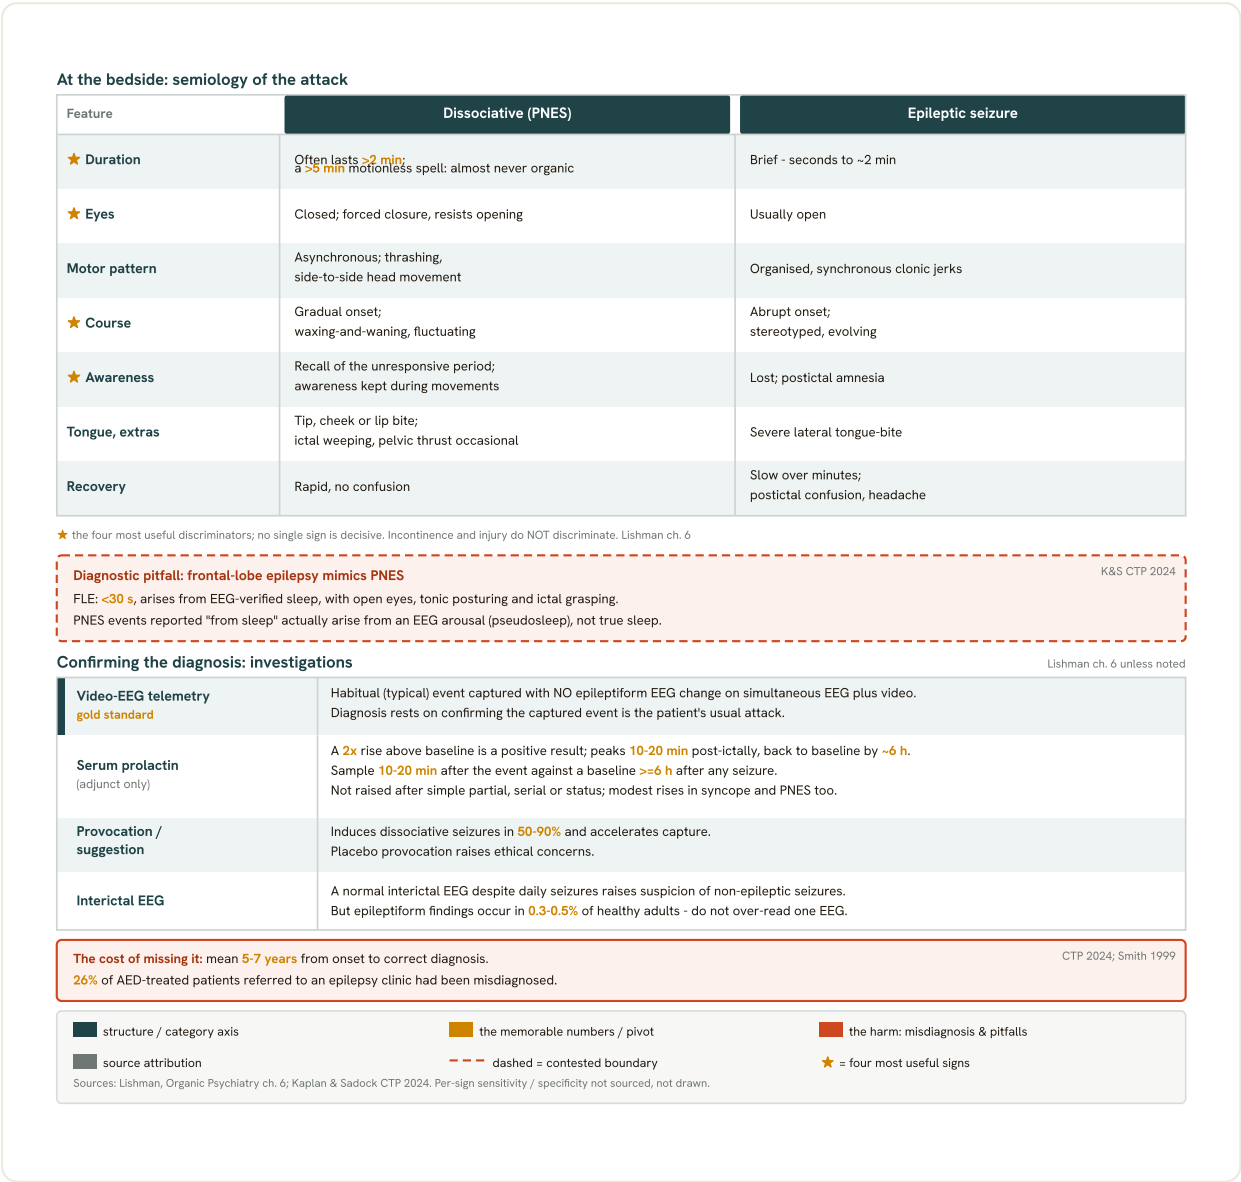


Fig 34.5 **PNES against epileptic seizures.** The four most useful bedside discriminators are long duration, eye closure, a fluctuating course and recall following unresponsiveness, yet no single sign is decisive. Video-EEG telemetry that captures the habitual event without epileptiform change is the gold standard, with serum prolactin and suggestion as adjuncts.

34.1 The governing principle: probability, not proof

Every sign shifts the odds; none closes the case. The single most important idea to carry into any answer on this topic is that no individual symptom or sign discriminates dissociative from epileptic seizures with certainty. Each feature behaves like a likelihood ratio, not a verdict. Epileptic seizures are electrically and clinically diverse, and frontal attacks in particular can be brief, bizarre and bilaterally motor while barely disturbing the scalp trace, so any single textbook feature of a dissociative attack can turn up in genuine epilepsy, and the occasional atypical dissociative event borrows the look of a convulsion. The clinician therefore reasons the way a

good diagnostician always does, by accumulating features that point the same way until the constellation becomes compelling, and by weighting what is observed against the reference standard rather than trusting a favourite sign. According to Lishman's Organic Psychiatry the most useful discriminators cluster into a small, memorable group, and it is their co-occurrence in one attack, corroborated by a reliable eyewitness or by a recording, that carries diagnostic weight. This is why a careful account from someone who watched the whole event, ideally more than one observer, is worth more than any laboratory result short of a captured attack.

RULE **Diagnose on the pattern, not on the single sign.** Because no one feature is fully discriminating, a confident bedside diagnosis rests on several concordant features in the same attack plus a trustworthy witness account. A resident who stakes the answer on one dramatic sign has missed the whole method, and will meet the counter-example in the viva.

34.2 Duration, onset and the tempo of the attack

Long, gradual and fluctuating all point away from epilepsy. Epileptic seizures are, as a class, brief and stereotyped: a convulsion runs its course in a couple of minutes and a focal seizure is usually shorter. A dissociative attack tends to last longer, and an event that continues **beyond two minutes** is common in the dissociative form and distinctly unusual in an epileptic one. Tempo matters as much as total length. Epileptic events begin abruptly and evolve through a fixed sequence, whereas a dissociative attack more often shows a **gradual onset** and a fluctuating, waxing-and-waning course, the motor activity rising, subsiding and rising again over the episode rather than marching through the ordered progression of a generalised fit. The most striking version of this is the patient found lying motionless and unresponsive for a long stretch who then recovers intact: a period of motionless unresponsiveness lasting **more than five minutes**, from which the person comes back without residual deficit, almost never has an organic basis. The reasoning behind these signs is mechanistic. An epileptic seizure is the behavioural expression of a self-limiting electrical storm that the brain terminates through its own inhibitory machinery within a tight window; a dissociative attack is not bound by that physiology, so it can be prolonged, can wax and wane with the patient's state and the surroundings, and can build up gradually rather than switch on.

34.3 The character of the movements

Disorganised and thrashing versus organised and rhythmic. When there is prominent motor activity, its character is one of the richest sources of discrimination. Epileptic convulsive movements are organised: bilateral clonic jerking is coarse, rhythmic and roughly synchronous, and it slows in frequency as the seizure exhausts itself. Dissociative motor activity tends instead to be **asynchronous**, with the limbs moving out of phase, and to take the form of thrashing, side-to-side head movement and violent flailing rather than the patterned jerking of a true convulsion. Two further movements are worth naming because they are classically taught even though they are neither common nor decisive: **pelvic thrusting** and truncal arching into **opisthotonus**, the arched posture sometimes called arc de cercle. These occur occasionally in dissociative attacks and are rare in epileptic ones, but they are not pathognomonic, and a candidate who calls pelvic thrusting proof of a functional attack will be caught out, because it can

arise in frontal seizures. The underlying logic again helps the memory. Synchronous, rhythmic, decrementing clonus reflects a single electrical rhythm driving both sides of the body from a common generator; asynchronous, variable, shifting movement has no such unifying driver, which is exactly why it looks disorganised and why its tempo can fluctuate through the attack.

34.4 The eyes and the face

Closed, resisting eyes and weeping favour a dissociative attack. The state of the eyes is one of the most reliable and easily observed signs. In the great majority of epileptic seizures the eyes are open, often deviated or staring; in dissociative attacks they are characteristically shut, and the finding that carries real weight is **forced eye closure** with active resistance when an examiner tries to prise the lids apart. That resistance implies a degree of motor control and engagement with the environment that a genuine ictal state does not permit, which is why it points so firmly toward a dissociative mechanism. Emotional expression during the attack tells a similar story. **Ictal weeping**, crying or distressed vocalisation during the event, occurs occasionally in dissociative attacks and is very rare in epileptic seizures, again because sustained, contextually coherent emotional behaviour is hard to reconcile with a generalised discharge that suspends organised cortical function. Neither sign stands alone, but eye closure in particular earns its place among the highest-yield discriminators and should be sought deliberately, and documented, in any witnessed or recorded attack rather than left to chance recall.

34.5 Awareness, responsiveness and recall

Being present during a supposedly generalised attack is the giveaway. Perhaps the most conceptually important discriminator concerns whether the patient was genuinely absent. In a generalised convulsive seizure consciousness is lost and the event is not remembered. In dissociative attacks patients commonly retain **preserved awareness** during bilateral motor activity and can afterwards recall the period in which they appeared unresponsive, and some respond to a bystander's intervention while the attack is running. Retained recall of a spell of apparent unresponsiveness, and any evidence that the patient was registering and reacting to the surroundings while both sides of the body were moving, are strong pointers, because bilateral motor activity with intact awareness is physiologically incongruent with a generalised epileptic discharge, which by definition disrupts consciousness. This is where a skilled history pays off. Gently asking the patient what they experienced during the attack, and asking witnesses whether the person seemed to respond to them or to a touch, extracts information that no single physical sign provides. It is also the feature most often overlooked, because clinicians assume that a dramatic motor event must have abolished consciousness and so never think to ask about it at all.

34.6 How the attack ends

Quick and clear-headed versus slow and confused. The recovery phase is as informative as the attack. After a generalised epileptic seizure the patient typically surfaces slowly over several minutes through a period of **postictal confusion**, disoriented, drowsy and not fully accessible. A dissociative attack characteristically ends with rapid, full recovery and a return to clear consciousness without that clouded aftermath. The contrast reflects the biology once more: an epileptic discharge leaves the cortex temporarily exhausted and inhibited, and that neuronal

aftermath takes time to clear, whereas a dissociative attack imposes no such metabolic debt and can therefore switch off cleanly. One important caution belongs here. Postictal tiredness and headache do not discriminate, because they are common after both kinds of attack and are driven as much by the physical exertion and the emotional aftermath as by any specific mechanism. It is the presence or absence of genuine confusion, the clouding of consciousness rather than mere fatigue, that carries the diagnostic signal, and conflating the two is a common way to misread the recovery phase and lose the point.

34.7 The features that do not help

Tongue-biting, incontinence and injury are widely over-trusted. Some signs enjoy a folk reputation as proof of epilepsy that the evidence does not support, and knowing their real value is a favourite examination point. Tongue-biting is the instructive case. Severe biting of the side of the tongue, **lateral tongue-biting**, is genuinely useful, being common in epileptic seizures and very rare in dissociative ones, but its worth is specific to that location: in dissociative attacks any bite is more likely to fall on the tip of the tongue, the cheek or the lip, so a report that the patient bit the tongue, without asking where, is nearly worthless. **Incontinence** of urine and physical injury are even less reliable. Incontinence is occasional in dissociative attacks and common in epilepsy but occurs in both, and injuries, including those from falls, are sustained in dissociative as well as epileptic events. Because these signs feel dramatic and consequential, they are habitually over-weighted, and treating them as clinching evidence for epilepsy is one of the surest routes to the wrong diagnosis.

TRAP **Taking incontinence, injury or a bitten tongue as proof of epilepsy.** These are the features candidates and clinicians most often over-read. Incontinence and injury occur in both kinds of attack, and only severe biting of the side of the tongue carries weight; a bite to the tip or the cheek does not. Marks, and diagnoses, are lost by presenting any of them as decisive.

34.8 The mimic that catches everyone: frontal lobe seizures

The dangerous error is the opposite one, calling a frontal seizure functional. The discriminators above are built around the generalised convulsion, and against it they work well. The great trap of the whole field is **frontal lobe epilepsy**, whose attacks are the principal epileptic imitators of dissociative seizures and the commonest reason a genuine seizure is wrongly labelled psychogenic. Frontal seizures can be brief, bizarre, hypermotor and emotionally florid, with thrashing and pelvic movements that look functional, and, crucially, they may produce little or no change on the scalp electroencephalogram, so even a recorded attack can mislead the unwary. Kaplan and Sadock's Comprehensive Textbook of Psychiatry sets out the features that pull the diagnosis back toward frontal epilepsy:

- a very short duration, often **under thirty seconds**, against the typically longer dissociative attack;
- onset directly out of electroencephalographically verified sleep, which dissociative attacks do not show;
- open rather than closed eyes during the event;

- tonic posturing of the limbs;
- ictal grasping and other stereotyped automatisms.

The lesson is one of humility. A stereotyped, very brief, nocturnal attack arising from confirmed sleep, with tonic posturing and open eyes, should raise frontal lobe epilepsy however functional the flailing looks, and this is precisely the situation in which the definitive investigation, prolonged video-EEG monitoring covered in its own topic, earns its keep. The wider behavioural and psychiatric presentation of frontal lobe epilepsy is developed separately; here it matters only as the great differential to hold in mind before committing to a functional label.

PITFALL

Brief, stereotyped, night-time hypermotor attacks with thrashing and pelvic movements are the ones most often dismissed as functional, and withdrawing antiepileptic drugs on that basis can be dangerous. When attacks are very short, arise from confirmed sleep and carry tonic posturing, treat frontal lobe epilepsy as the leading possibility and seek monitoring before you relabel the diagnosis.

34.9 The highest-yield cluster and reading the whole attack

Assemble the features and let them vote. Pulled together, the discriminators that Lishman's Organic Psychiatry rates as most useful form a compact set: a long attack duration, eye closure, fluctuating motor activity, and recall of the period of unresponsiveness. When several of these appear in one attack, corroborated by a good witness, the probability of a dissociative mechanism becomes high; when they are absent and the event is a brief, organised, amnesic convulsion with slow confused recovery, epilepsy is far more likely. The accompanying table lays the individual features side by side so that the whole phenotype, rather than any one line of it, can be weighed at a glance.

MNEMONIC

FRED

The four features that shift the odds hardest toward a dissociative attack: **F**luctuating, waxing-and-waning motor activity; **R**ecall of the unresponsive period; **E**ye closure; and a long **D**uration. The point of grouping them is that they carry weight together, not one at a time, and each is a discriminator this topic actually states.

FEATURE	FAVOURS DISSOCIATIVE (PNES) ATTACK	FAVOURS EPILEPTIC SEIZURE
Attack duration	Often prolonged	Brief, self-limited
Onset and course	Gradual, waxing and waning	Abrupt, stereotyped sequence
Prolonged motionless spell with full recovery	Suggestive, rarely organic	Distinctly unusual
Motor pattern	Asynchronous, thrashing, side-to-side head movement	Organised, roughly synchronous clonus

FEATURE	FAVOURS DISSOCIATIVE (PNES) ATTACK	FAVOURS EPILEPTIC SEIZURE
Pelvic thrusting, opisthotonus	Occasional, not pathognomonic	Rare
Eyes	Closed, often forcibly resisting opening	Usually open
Ictal weeping	Occasional	Very rare
Awareness and recall	May be preserved, with later recall	Lost, not remembered
Recovery	Rapid, clear-headed	Slow, with postictal confusion
Tongue-biting	Tip, cheek or lip, if any	Severe, to the side
Incontinence and injury	Non-specific, can occur	Non-specific, can occur

A feature-by-feature comparison of dissociative and epileptic attacks; the diagnosis is read from the balance of the whole column, since no single row is decisive on its own.

over 2 min	over 5 min	under 30 s
ATTACK DURATION COMMON IN PNES	MOTIONLESS YET RECOVERS: RARELY ORGANIC	FRONTAL LOBE SEIZURE, THE MIMIC

IN INDIA

In much of Indian practice the diagnosis has to be made on semiology alone, because prolonged video-EEG monitoring is concentrated in a few tertiary centres and is out of reach for most patients. This raises the premium on a meticulous eyewitness history, increasingly from a smartphone recording of the attack that families can be coached to capture. The older label of hysteria still colours how these presentations are received, and the stigma attached both to epilepsy and to a psychiatric attribution means patients frequently arrive already committed to antiepileptic drugs and reluctant to accept a functional explanation. Getting the clinical differentiation right, gently and alongside the family, matters most exactly where confirmatory testing is scarce.

GOOD TO REMEMBER

When you cannot get an attack on video, chase the account instead: how long it lasted, whether the eyes were shut and resisted opening, whether the movements were rhythmic or thrashing, whether the person remembers the spell or responded to anyone during it, and how quickly they came round afterwards. Ask a witness to record the next event on a phone. Several concordant soft signs beat one dramatic sign every time.

No single sign is decisive: the diagnosis of a dissociative attack rests on a cluster of long duration, eye closure, fluctuating asynchronous movement, preserved awareness with later recall and a rapid clear recovery, while very brief nocturnal attacks with tonic posturing must always raise frontal lobe epilepsy as the great mimic.

35 · Video-EEG, serum prolactin and other investigations in PNES diagnosis

The whole weight of a confident diagnosis here falls on investigation, because the bedside is never quite enough. The features of the attack itself, set out in the companion topic on the clinical features that differentiate the two events, shift the odds but never close the case, and both errors that follow from stopping at the bedside are costly: a person committed to antiepileptic drugs for an illness they do not have, or a genuine seizure waved away as functional. What settles the question is a recorded event, and the marks in this area cluster around a single disciplined idea, that the diagnosis is made positively by capturing the attack, not negatively by a normal routine test. This topic walks the investigative pathway in the order a clinician actually uses it: the reference standard of simultaneous video and electroencephalography, the interictal trace and its two opposite traps, serum prolactin and its narrow window, provocation, the sleep-onset discriminator, and the sobering delay that makes getting all of this right a matter of years.

What a dissociative attack actually is, and the shifting terminology it has travelled under, belongs to its own companion topic and, more broadly, to the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes. The general fundamentals of the electroencephalogram, and the distinction between ictal and interictal discharges, are developed in their own right and are used here only as they bear on this one diagnostic question. Delivering the diagnosis to the patient and its stepped management, the psychiatric comorbidities and predisposing factors, and the genuinely difficult patient who has both dissociative and epileptic seizures at once are each treated separately. Here we stay with the tests: what each one can prove, what it cannot, and how they are sequenced.

35.1 The reference standard: ictal video-EEG telemetry

Capture the attack with its brain trace and the argument is over. The single investigation that settles the diagnosis is **video-EEG telemetry**, prolonged simultaneous recording of the electroencephalogram alongside time-locked video of the patient, and per Lishman's Organic Psychiatry it is the accepted gold standard for the condition. Its power lies in recording the clinical event and the brain's electrical activity at the same instant, so that semiology and cortical rhythm can be read against each other rather than inferred separately. The diagnostic move it permits is a positive one. According to Kaplan and Sadock's Comprehensive Textbook of Psychiatry, the diagnosis is established when the patient's habitual attack is captured and the trace shows no accompanying epileptiform change during the event: the typical spell unfolds on camera while the electroencephalogram stays free of any ictal discharge. Two conditions have to hold for this to mean anything. The recorded event must be confirmed by the patient or a relative as their usual attack and not some lesser or atypical spell, because a normal trace during an event nobody recognises proves nothing about the attacks that matter. And the trace reflects cortical seizures visible at the scalp, so deep or very small discharges can escape surface electrodes, which is why the method leans on capturing the full, typical semiology rather than a fragment. Held to those conditions, a captured habitual attack with a clean ictal trace is as close to certainty as this field offers.

RULE **The diagnosis is positive, not an absence.** Non-epileptic attacks are confirmed by recording the patient's own typical event and showing that no ictal electrographic change accompanies it, with the event verified as habitual. A normal test in the gaps between attacks is not the same thing and never carries the same weight.

35.2 The interictal EEG: two opposite traps

A normal trace does not rule epilepsy out, and an abnormal one does not rule it in. Between attacks the electroencephalogram is far weaker evidence than most people assume, and it misleads in both directions. On one side, a completely normal **interictal EEG** in a patient reporting seizures every day is genuinely suggestive of non-epileptic attacks, because uncontrolled daily epilepsy usually leaves some interictal footprint. Yet that inference is soft, not hard, because roughly one-third of people with epilepsy show no interictal abnormality at all when awake, so a clean trace is entirely compatible with real epilepsy and can never exclude it. On the other side lies the more dangerous error. Interictal and even frankly epileptiform discharges can appear in people whose only attacks are dissociative, and epileptiform findings turn up in about **0.3 to 0.5 per cent** of entirely healthy adults, so a sharp wave on the trace is not a licence to diagnose epilepsy and start drugs. Lishman's Organic Psychiatry is explicit that the interictal record has to be read in the light of the history and the semiology and never in isolation. The practical upshot is that the routine EEG, ordered and interpreted on its own, is one of the commonest sources of diagnostic error in this whole area, capable of manufacturing a false epilepsy label out of a benign blip or false reassurance out of a normal study. It is a supporting tool, not an arbiter.

TRAP **Reading a normal interictal EEG as proof of non-epileptic attacks, or a minor epileptiform blip as proof of epilepsy.** Both are examination favourites. A substantial minority of people with genuine epilepsy have normal traces between seizures, and a small fraction of healthy people carry epileptiform discharges, so the interictal study confirms neither diagnosis. Candidates who let it decide the case lose the mark.

35.3 Serum prolactin: the physiology and its narrow window

A generalised discharge that reaches the hypothalamus drives prolactin up, but only briefly. The rationale for the **serum prolactin** test is that a generalised epileptic discharge spreading through hypothalamic pathways transiently releases prolactin from the anterior pituitary, a hormonal signature that a purely psychological event should not produce. The kinetics are the whole test and must be known precisely. Following a tonic-clonic seizure the serum level climbs and reaches its peak at about **10 to 20 minutes**, then falls back, returning to baseline after roughly 6 hours, as characterised by Lishman's Organic Psychiatry. Everything that makes the test useful and everything that makes it fail follows from that curve. Because the rise is short-lived, a sample drawn too late catches a level already on its way down or back to normal, which is then read, wrongly, as a negative result. Because the release depends on a discharge large enough to engage the hypothalamus, seizures that stay small or focal may never move the hormone at all. And because prolactin also responds to physiological stress, the specificity of any single raised

value is limited. The accompanying figure sketches this brief postictal rise and its return; the sampling rules that follow simply try to catch the hormone at the top of that curve.

35.4 Sampling prolactin correctly and reading the result

One value proves nothing; the test is a timed pair. A prolactin result is interpretable only as the relationship between a well-timed postictal sample and the patient's own baseline, and getting either wrong wastes the test. Lishman's Organic Psychiatry sets out the protocol:

- draw the postictal sample at the top of the curve, within that early peak window and not later;
- take a paired baseline sample at least 6 hours after any seizure, when the hormone has certainly returned to rest;
- match that baseline to approximately the same time of day, because prolactin has its own diurnal rhythm.

Only against that baseline does the postictal figure mean anything, and the threshold conventionally taken as a significant, positive result is a rise to **at least twice the baseline** value, although the exact cut-off has varied between studies and is not a single agreed number. A doubling or more, sampled at the right moment against a proper resting level, supports a preceding generalised or complex partial seizure; a value below that threshold is informative only if you are confident the sample was drawn inside the window and the event was one that should have raised the hormone.

> **GOOD TO REMEMBER**
>
> If you are going to send a prolactin at all, send it properly or not at all: one tube at the peak of the rise, as soon as possible after the attack, and a second, paired baseline tube drawn well clear of any seizure and matched for time of day. A lone postictal value with no baseline, or one drawn too late, is uninterpretable, and an uninterpretable test is worse than no test because it invites a false conclusion.

35.5 Why prolactin is an adjunct, not a rule-out test

Too many epileptic seizures fail to raise it, and too many other events do. The limitations of the test are as examinable as its rules, because they define exactly what a normal or a raised value is allowed to mean. Per Lishman's Organic Psychiatry, serum prolactin is not reliably raised after several kinds of genuinely epileptic activity, and it rises in some events that are not epileptic at all. On the false-negative side, it is typically not elevated after simple partial seizures, after serial seizures and after **status epilepticus**, and it is a less dependable marker after complex partial seizures than after a full tonic-clonic convulsion. On the false-positive side, modest rises occur in **syncope** and in dissociative attacks themselves. So a failure to rise never proves an event was non-epileptic, and a modest rise never proves it was epileptic.

MNEMONIC

Simple, Serial, Status

The three epileptic situations after which serum prolactin is characteristically *not* raised, so a normal value cannot be trusted to exclude a seizure: **S**imple partial seizures, **S**erial seizures and **S**tatus epilepticus, with the complex partial seizure a fourth, less reliable, case. Each is a situation this source names explicitly.

PITFALL

The classic clinical mistake is to read a normal postictal prolactin as evidence that an attack was psychological, when in fact it followed a simple partial or complex partial seizure, or was drawn outside the window. The mirror-image error is to treat a modest rise as proof of epilepsy when syncope, or even a dissociative attack, can produce one. Prolactin helps sort a probable tonic-clonic seizure from a probable functional event; it does nothing else safely.

Where the test still earns its keep is captured by the formal verdict on it. According to Lishman's Organic Psychiatry, the **American Academy of Neurology** judged postictal serum prolactin a useful adjunct for separating generalised and complex partial seizures from psychogenic events, while insisting that it is an adjunct and not a stand-alone test. With prolonged video-EEG now more widely available the assay has fallen out of routine favour, but a negative result after an apparently convulsive attack retains real practical value, and in settings without ready telemetry that modest role is not to be dismissed.

35.6 Provocation and induction: bringing an attack on

If attacks are too infrequent to catch, you can try to precipitate one. Prolonged telemetry only works if an event obligingly occurs while the patient is wired up, and many people do not attack often enough for that. The response is **provocation**, the deliberate use of induction to precipitate a habitual attack under recording so that its nature can be settled on the trace. Per Lishman's Organic Psychiatry, the non-specific activation procedures already built into a routine study, **hyperventilation** and **photic stimulation**, together with verbal suggestion, will induce a dissociative attack in a large proportion of susceptible patients, quoted across studies as somewhere between **50 to 90 per cent**, and they shorten the time to a captured event. Simple suggestion alone, without any elaborate placebo, provoked a typical event in about two-thirds of patients in one series and cut the need for prolonged telemetry in nearly half. The obvious efficiency of this has to be weighed against its ethics. Techniques that rely on active deception, for instance injecting saline described as a drug that will bring on a fit, raise real problems of consent and of the very trust on which the subsequent conversation about a functional diagnosis depends, and many units now prefer transparent activation and honest suggestion to any frank placebo. However it is elicited, the captured event still has to be confirmed as the patient's usual attack to count.

35.7 Pseudosleep and the sleep-onset discriminator

An attack that seems to come out of sleep may be coming out of wakefulness. One elegant discriminator emerges only because video and electroencephalography run together, and it turns on what "sleep" actually means on the trace. Patients with dissociative attacks not uncommonly report that their events arise from sleep, which sounds reassuringly organic, but the recording tells a different story. According to Lishman's Organic Psychiatry, these attacks in fact begin from a state of **pseudosleep**: the patient appears to be asleep, yet the trace shows an arousal to wakefulness in the moments before the attack starts, so the event is really arising from the waking brain behind closed eyes. This matters because it is precisely the feature that the main epileptic mimic does not share. Frontal lobe epilepsy, the seizure disorder that most often imitates dissociative attacks, can genuinely arise directly out of electrographically documented sleep, so an attack that emerges from a verified sleep trace points toward epilepsy while one that emerges from pseudosleep points away from it. The wider behavioural and psychiatric presentation of frontal lobe epilepsy is developed in its own right; here it matters only as the organic counterpoint that makes the pseudosleep sign useful, and the finding is available only when the attack is caught with simultaneous video and electroencephalography.

35.8 Sequencing the tests, and the cost of getting it wrong

The tests form a hierarchy, and the reason to climb it quickly is measured in years. No single investigation is used in isolation; they form a ladder, and the accompanying table lays out what each rung can and cannot prove. Beyond these, neuroimaging and neuropsychological assessment have supporting roles in the wider work-up without confirming or excluding the diagnosis, and are named here only to be placed. The stakes in getting up this ladder are high, because the natural history of a missed diagnosis is long. According to Kaplan and Sadock's Comprehensive Textbook of Psychiatry the mean latency from the first manifestation to the correct diagnosis runs to **5 to 7 years**; Lishman's Organic Psychiatry records that such events are commonly treated as epileptic for more than 3 years, and that among antiepileptic-drug-treated patients referred on to a specialist epilepsy service, 26 per cent had in fact been misdiagnosed. Those years are not benign: unnecessary drugs with their own adverse effects, psychological treatment not given, and a diagnostic identity that hardens the longer it stands. That is the real argument for reaching the reference standard early rather than accumulating inconclusive routine tests.

INVESTIGATION	WHAT A POSITIVE OR TYPICAL RESULT SHOWS	MAIN LIMITATION
Ictal video-EEG telemetry	Habitual attack captured with no ictal EEG change: the positive diagnosis	Needs an event during recording; deep discharges can escape scalp electrodes
Interictal (routine) EEG	A normal trace despite frequent daily attacks is suggestive of non-epileptic events	Normal in much genuine epilepsy; epileptiform blips occur in healthy people
Serum prolactin (paired, timed)	A sufficient postictal rise supports a generalised or complex partial seizure	Not raised after several seizure types; rises in syncope and dissociative attacks
Provocation and suggestion	Precipitates a habitual attack under recording, accelerating capture	Deceptive placebo raises consent and trust concerns; event must be confirmed habitual
Video-EEG in apparent sleep	Pseudosleep arousal before onset points away from epilepsy	Requires simultaneous video and EEG; interpretation is specialist

The investigative ladder in psychogenic non-epileptic attacks, from the confirmatory reference standard down to the supporting adjuncts; no rung below the captured attack is diagnostic on its own.

10 to 20 min	**at least 2x**	**50 to 90%**	**5 to 7 years**
POSTICTAL PROLACTIN PEAK	RISE COUNTED AS POSITIVE	ATTACKS PROVOKED BY SUGGESTION	TYPICAL DELAY TO DIAGNOSIS

IN INDIA

Prolonged video-EEG telemetry sits in a handful of tertiary centres, so across much of the country the diagnosis has to lean on the cheaper, more available adjuncts, and this is exactly where they are most often misused. A single unpaired serum prolactin, drawn whenever the patient happens to reach the laboratory, is ordered widely and then interpreted as if a normal value excluded a seizure, which it cannot. The disciplined answer where telemetry is scarce is a properly timed, paired prolactin used only to support a probable convulsion, a good home or smartphone recording of a typical attack, and pragmatic activation during a routine study, with frankly deceptive placebo methods best avoided given how heavily the later therapeutic relationship depends on the patient's trust.

The diagnosis of non-epileptic attacks is made positively, by capturing the patient's own habitual event on video-EEG with no accompanying ictal discharge; the interictal EEG, serum prolactin and provocation only support that finding, no normal result ever excludes epilepsy, and a properly timed twofold prolactin rise means a probable generalised seizure and nothing more.

36 · Psychiatric comorbidity and predisposing factors in PNES

Vulnerability is the frame here, not causation. A psychogenic non-epileptic seizure does not arise from nowhere; it grows on a psychiatric and developmental substrate, and the marks in this topic sit on knowing what that substrate is, who carries it and how it is assembled into a

formulation. Examiners return to two things again and again: the comorbidity profile that separates these patients from those with epilepsy, and the discipline that no amount of comorbidity, however striking, actually establishes the diagnosis. A candidate who can list the predisposing, precipitating and maintaining factors, weight them correctly, and still say that the diagnosis rests elsewhere has understood the whole topic. The clinical stakes are the same as the examination ones, because a formulation of vulnerability is what turns a confirmed diagnosis into a treatment plan the patient can accept.

This account stays with the who and the why. What a **dissociative seizure** actually is, the terminology it travels under and its place within the family of functional neurological disorders belong to the companion topic and, for the general construct, to the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes. How these attacks are told apart from epileptic seizures at the bedside, how they are confirmed with video-electroencephalographic capture and serum prolactin, how the diagnosis is communicated and managed in stepped fashion, and the genuinely difficult patient who has both dissociative and epileptic seizures at once, are each developed in their own right. Here the material is narrower: the epidemiology of the vulnerable patient, the psychiatric disorders that cluster around these attacks, the childhood and illness-behaviour history that predisposes, and the biopsychosocial architecture that holds it all together.

36.1 The vulnerable patient: who develops dissociative seizures

Young, disproportionately female, and sitting in the epilepsy clinic already. The demographic signal is consistent enough to be examinable. According to Kaplan and Sadock's Comprehensive Textbook of Psychiatry, dissociative seizures affect women in their reproductive years at approximately a **3:1 female-to-male** ratio, are most frequent in young adults, and yet occur across the whole lifespan, so neither an older nor a male patient can be excluded on demographic grounds alone. The second half of the epidemiology matters more for daily practice, because it locates these patients where the neurologist and the psychiatrist actually meet them. Far from being a rarity, dissociative seizures account for up to **30 per cent** of admissions to specialist epilepsy monitoring units, and about one in five of the patients referred for evaluation of apparently intractable epilepsy prove, on capture, to have dissociative seizures rather than a drug-resistant epilepsy. The lesson embedded in those figures is that the population at risk is not a separate psychiatric caseload but a large fraction of the very patients being worked up as difficult epilepsy. A young woman whose seizures are not responding to escalating antiseizure medication is, on base rates alone, a patient in whom the possibility should already be alive, and the reproductive-age preponderance is why the disorder so often intersects with pregnancy and the social pressures around young women.

36.2 The comorbidity profile that separates PNES from epilepsy

Trauma, PTSD and personality disorder discriminate; depression does not. The single most testable idea in the comorbidity of these attacks is a comparative one, and it is easy to get half-right and lose the mark. Both patients with dissociative seizures and patients with epilepsy carry a heavy psychiatric burden, so the discriminating question is not whether a disorder is present but whether it is present more often than in epilepsy. Kaplan and Sadock's Comprehensive

Textbook of Psychiatry draws the line cleanly: prior traumatic experiences, **post-traumatic stress disorder** and **personality disorder** are more frequent in dissociative seizures than in epilepsy, whereas **mood disorder** is frequent in dissociative seizures but not significantly more so than in epilepsy. The direction of that last clause is the whole examination point, and reversing it is the classic slip. Depression is common in both groups because chronic, frightening, socially disabling paroxysmal disorders are depressing whatever their mechanism, so a high rate of low mood does nothing to separate the two. It is the trauma-spectrum burden, the post-traumatic and personality pathology, that is genuinely over-represented in the dissociative group, which is why the trauma history is worth taking with particular care even though it will never, by itself, decide the case. The accompanying table lays the comparison out so that the direction of each row can be read at a glance.

PSYCHIATRIC DISORDER	IN DISSOCIATIVE SEIZURES, COMPARED WITH EPILEPSY
Prior traumatic experiences	More frequent than in epilepsy
Post-traumatic stress disorder	More frequent than in epilepsy
Personality disorder	More frequent than in epilepsy
Mood disorder (depression)	Frequent, but not significantly more than in epilepsy

The comparative comorbidity profile after Kaplan and Sadock's Comprehensive Textbook of Psychiatry; only the trauma-spectrum rows discriminate, and the mood-disorder row is the one candidates most often invert.

36.3 Childhood adversity and abuse as the developmental substrate

The recurring theme is early adversity, above all abuse. Behind the adult comorbidity lies a developmental story that Lishman's Organic Psychiatry treats as a common underlying theme rather than a universal law. Adverse and traumatic childhood experiences run through these histories, with a particularly highlighted association with **childhood sexual, physical and emotional abuse**. The mechanistic reading, which the exam rewards over rote recall, is that early abuse and adversity shape the very capacities that dissociative seizures exploit: an ingrained tendency to dissociate under threat, a body that has learned to express distress somatically, and an attachment and coping style forged in an environment where direct expression of need was unsafe or futile. What the candidate must hold alongside this, though, is the caveat the source itself insists on. This developmental adversity is not unique to dissociative seizures; it is shared with the other somatoform and dissociative disorders, so a childhood abuse history raises the prior probability of a functional mechanism in general without pointing specifically at seizures and without proving anything on its own. The two errors to avoid are opposite in direction. One is to treat a disclosed abuse history as confirmatory; the other, subtler and commoner, is to treat its absence as exclusionary, forgetting both that adversity is a theme and not a criterion and that abuse is chronically under-disclosed, especially at a first meeting and where cultural shame is heavy.

PITFALL

Do not read a negative trauma history as evidence against the diagnosis, and do not go hunting aggressively for buried abuse to shore one up. Childhood adversity is a common theme, not a necessary ingredient, and disclosure is often delayed by shame and by the very dissociative style that predisposes to the attacks. Pressing a reluctant patient to produce a trauma narrative early can rupture engagement and reads, wrongly, as making the diagnosis contingent on a confession.

36.4 A history of medically unexplained symptoms

The seizure is usually not the first functional presentation. One of the most useful pieces of history in these patients is not about seizures at all. Dissociative seizures tend to arrive in a person with an established pattern of somatic distress, and Lishman's Organic Psychiatry records that up to **80 per cent** of patients with these attacks have a history of previous **medically unexplained presentations**. The attack, in other words, is often the latest and most dramatic expression of a longer somatising trajectory rather than a bolt from a clear sky, and a careful review of past consultations frequently turns up earlier functional symptoms, repeated investigations that came to nothing, and a relationship with medical services already shaped by unexplained complaints. This matters clinically in two ways. It supports the formulation, because a lifelong tendency to somatise is the diathesis on which a dissociative seizure disorder is built, and it warns against the reflex of investigating each new symptom afresh as though the slate were blank, which entrenches illness behaviour and risks iatrogenic harm. It also softens the disclosure that comes later, because a patient who can be shown the pattern of their own history is often more able to accept a functional account than one confronted with the seizures in isolation. The figure is worth carrying, but it describes a tendency and a prior probability, not a gate through which every patient must pass.

36.5 The biopsychosocial formulation: predisposing, precipitating, maintaining

Do not list the factors; assemble them across three time-frames and three domains. The mature way to hold everything above is the formulation that Lishman's Organic Psychiatry sets out, in which vulnerability is organised along two axes at once: a temporal axis of predisposing, precipitating and maintaining factors, and a domain axis of biological, psychological and social contributions. The predisposing factors are the standing vulnerabilities that make a person the kind of person in whom these attacks can develop, spanning all three domains: a biological substrate such as **learning difficulty**; the psychological traits of somatising and dissociative tendency, **avoidant coping** and an underlying personality or mood disorder; and the social conditions of poor family functioning and the modelling of attacks on others with epilepsy in the patient's environment. The precipitating factors are what tips a vulnerable person into the first attacks, and Lishman's Organic Psychiatry names adverse life events and acute episodes of panic or syncope here. The maintaining factors are what keeps the attacks going once begun, and they are where treatment often has most leverage: the benefits of the **sick role**, agoraphobic avoidance and safety behaviour, and the confused or anxious reactions of the patient and of carers to the diagnosis itself. Reading them together shows why the disorder is self-perpetuating: the

precipitant lights the fire, and the sick-role and avoidance keep feeding it long after the original trigger has passed.

MNEMONIC

Three Ps, three domains

Formulate every case on a grid. Down the side, the three Ps: **P**redisposing (standing vulnerability), **P**recipitating (what triggered the first attacks) and **P**erpetuating or maintaining (what keeps them going). Across the top, the three domains: biological, psychological and social. The predisposing row fills all three domains; the value of the grid is that it forces you to place each factor in both its time-frame and its domain rather than reciting an undifferentiated list.

RULE **Vulnerability is formulated, not merely listed.** The examinable move is to sort every factor into predisposing, precipitating or maintaining, and into biological, psychological or social. A flat list of risk factors earns little; a formulation that separates the standing diathesis from the trigger and from the perpetuating loop shows the reasoning and tells you which factors are actually modifiable.

TIME-FRAME	DOMAIN	FACTORS (AFTER LISHMAN'S ORGANIC PSYCHIATRY)
Predisposing	Biological	Learning difficulty
Predisposing	Psychological	Somatising and dissociative traits; avoidant coping; personality or mood disorder
Predisposing	Social	Poor family functioning; modelling of attacks on others with epilepsy
Precipitating	Psychological or physical	Adverse life events; an acute episode of panic or syncope
Maintaining	Psychological	Benefits of the sick role; agoraphobic avoidance and safety behaviour
Maintaining	Social	Confused or anxious reactions of the patient and of carers to the diagnosis

The predisposing, precipitating and maintaining formulation; the standing vulnerability spreads across all three domains, while the trigger and the perpetuating loop are where the timeline earns its keep.

36.6 Pre-ictal autonomic arousal and the panic overlap

Most attacks are preceded by a surge of autonomic symptoms. One phenomenological feature earns its place among the predisposing and precipitating material because it links the physiology of arousal to the attack itself. Lishman's Organic Psychiatry records that about **two-thirds** of patients with dissociative seizures report symptoms of **autonomic arousal** preceding their seizures, a pre-ictal build-up that typically includes:

- hyperventilation;

- palpitations;
- paraesthesiae;
- a dry mouth;
- perspiration.

The clinical significance of this cluster is that it overlaps almost completely with a panic attack and with presyncope, which is exactly why acute panic and syncope appear among the precipitating factors and why the boundary with an anxiety disorder is so blurred. A person prone to panic who then dissociates as the arousal peaks has a plausible route from autonomic surge to attack, and for many patients the seizure functions as an escape from an intolerable state of arousal. This is where the topic touches, without re-teaching, the ictal fear and panic phenomena of epilepsy and the interictal anxiety disorders that are developed elsewhere; the point owned here is that the arousal is a common pre-ictal experience and a maintaining engine, not that it is a diagnostic sign. It also has a practical yield, because a patient who recognises the arousal prodrome can be taught to intervene at that moment, which is the entry point for several of the psychological treatments described in the management topic.

36.7 Comorbidity raises probability, it does not make the diagnosis

The psychiatric history is context, not proof. Everything in this topic is a statement about probability, and the cardinal error is to let it harden into a diagnosis. Kaplan and Sadock's Comprehensive Textbook of Psychiatry is explicit that clinicians should not assume a dissociative-seizure diagnosis on the basis of comorbid psychiatric disorder alone, such as post-traumatic stress disorder or a personality disorder. The reasoning is straightforward and worth being able to state. Trauma, post-traumatic stress disorder and personality pathology are common in the general population and commoner still among people with epilepsy, so a patient can carry every one of these and still have entirely epileptic seizures; conversely, a patient with none of them can have dissociative seizures, since the risk factors are enrichments of probability, not necessary conditions. The diagnosis is made where it has always been made, on the attack itself, from a captured event and its electrical correlate and the clinical semiology, all of which are developed in the semiology and investigations topics. The comorbidity and predisposing factors do real work in raising suspicion, in shaping the formulation and in guiding treatment, but they sit on the wrong side of the line to establish the diagnosis, and the most dangerous version of the error runs the other way, dismissing a patient's genuine seizures as functional because their psychiatric history is florid.

TRAP **Diagnosing dissociative seizures from the psychiatric history.** The classic error is to reason that a young woman with an abuse history, post-traumatic stress disorder and a personality disorder whose seizures resist medication must have functional attacks. Every clause raises the prior, and none of them is diagnostic. The examiner wants you to say that comorbidity informs suspicion but the diagnosis is made on the captured attack, and that reversing this has sent patients home with untreated epilepsy.

36.8 Formulating vulnerability in the individual patient

The purpose of all of this is a formulation the patient can accept. Pulled together, the material assembles into one working picture: a young, often female patient who carries a trauma-spectrum burden heavier than that seen in epilepsy, a developmental history in which adversity and abuse are a recurring theme, a long track record of medically unexplained symptoms, and attacks heralded by a surge of autonomic arousal that shades into panic. Laid on the predisposing, precipitating and maintaining grid, these stop being a list of risk factors and become a shared explanation of how this person came to have these attacks and what keeps them going, which is the account that makes the diagnosis credible and the treatment intelligible when the time comes to deliver it. The formulation is where comorbidity earns its clinical keep, and it is also where the discipline of the topic is honoured, because a formulation is offered as an understanding of vulnerability, never smuggled in as the evidence for the diagnosis.

3:1	**up to 30%**	**up to 80%**	**~two-thirds**
FEMALE-TO-MALE, REPRODUCTIVE YEARS	OF EPILEPSY MONITORING UNIT ADMISSIONS	WITH PRIOR UNEXPLAINED SYMPTOMS	REPORT PRE-ICTAL AUTONOMIC AROUSAL

IN INDIA

The reproductive-age, predominantly female profile lands on a specific social reality. Presentations still often reach services under the older idiom of hysteria, and the attacks frequently carry a heavy load of family and marital distress, dowry pressure and constrained agency for young women, which functions both as precipitant and as maintaining context. Childhood adversity and abuse are under-disclosed at the best of times and more so where shame and family honour are in play, so a negative first-visit history is especially weak evidence. The stigma attached both to epilepsy and to a psychiatric attribution means many patients arrive already committed to antiseizure drugs and to a physical explanation, and the formulation has to be built gently, often with the family in the room, if it is to be accepted at all.

GOOD TO REMEMBER

Take the trauma history and the illness-behaviour history with real care, because they build the formulation and soften the eventual disclosure, but write them into the record as vulnerability and context, never as the basis of the diagnosis. The two histories that most repay the effort are the map of past medically unexplained presentations and the account of the pre-ictal arousal, because both feed directly into treatment: the first reframes the pattern for the patient, and the second gives them a moment to intervene.

Trauma, post-traumatic stress disorder and personality disorder are over-represented in dissociative seizures while mood disorder is not, childhood adversity and a long history of medically unexplained symptoms are recurring predisposing themes, and the whole picture is formulated as predisposing, precipitating and maintaining factors across biological, psychological and social domains, yet none of it makes the diagnosis, which is read from the captured attack.

37 · Communicating the diagnosis and stepped psychological and pharmacological management of PNES

The consultation in which you tell the patient is itself the first dose of treatment. Once the investigations have settled that the attacks are **psychogenic non-epileptic seizures** and not epilepsy, the clinical problem changes shape. It stops being a diagnostic puzzle and becomes a therapeutic one, and the pivot on which everything after it turns is the conversation in which the diagnosis is handed over. Handled badly, that conversation loses the patient before treatment has begun, sometimes for good; handled well, it can begin to reduce the attacks on its own. The marks in this topic sit almost entirely in that transition: how to say what has been found, how to keep an easily alienated patient engaged, and how to build the stepped plan that follows the news.

This account picks the story up at the point of disclosure. How the diagnosis is reached and confirmed, including video-EEG and serum prolactin, belongs to the investigations topic; how these attacks are told apart from epileptic seizures at the bedside, the psychiatric comorbidity and predisposing factors that cluster around them, and where the disorder sits within the family of **functional neurological disorder** are each developed elsewhere, the general construct of conversion and functional neurological symptoms in the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes. What is owned here is narrower and more practical: the delivery of the news, the non-confrontational stance it demands, the safe withdrawal of unnecessary antiepileptic drugs, and the stepped, multidisciplinary treatment in which psychological therapy, not medication, does the definitive work.

37.1 Delivering the diagnosis is the pivotal step

Engagement is the treatment's first hurdle and its most fragile. The most important and often the most difficult task in the whole management of these patients is engaging them in treatment at all, and the vehicle for that engagement is the way the diagnosis is communicated. This is why the disclosure is treated as a management step in its own right rather than as a formality that precedes the real work. Lishman's Organic Psychiatry, drawing on the observations of Shen and Mellers, makes the strong version of the point: the **communication of the diagnosis** is potentially therapeutic in itself and can reduce the frequency and severity of the attacks, an effect that is most marked in recent-onset cases where the pattern has not yet consolidated. The mechanism is intuitive once stated. A patient who has been investigated repeatedly, often labelled with refractory epilepsy, and treated with escalating antiseizure drugs to no effect, has usually accumulated fear, uncertainty and a sense that something dangerous is being missed. A clear, credible account that names the condition, explains what the attacks are and are not, and offers a route out of the impasse removes much of the ambiguity that has been feeding the symptom. The clinical implication is direct: the encounter deserves unhurried time, a private setting and a prepared clinician, because it is doing more than conveying a result. It is beginning the intervention, and the early consultation is a window that closes as the disorder becomes entrenched.

37.2 The non-confrontational imperative

The stance governs the content. The single principle that governs the entire encounter is that it must be **non-confrontational**. Patients with these attacks are extremely sensitive to any suggestion, however oblique, that they are producing the events deliberately or putting them on, and the merest hint of that reading will end the therapeutic relationship. The vulnerability is not perverse; it is structural. In lay understanding the idea of a psychological seizure sits uncomfortably close to an accusation of faking, and a patient who hears the diagnosis as a charge of fabrication will feel disbelieved, dismissed and, quite reasonably, angry. Because the symptoms are involuntary and genuinely experienced, treating them as if they were feigned is both wrong and destructive. The practical consequence is that every element of the explanation has to be framed positively and explicitly around genuineness. The clinician states plainly that the attacks are real, that the patient is believed, and that the diagnosis is not a verdict of malingering. Only once that ground is secured does the rest of the conversation become possible, which is why the stance is logically prior to the content rather than a matter of bedside manner laid over it.

RULE **The delivery must be non-confrontational above all else.** Because these patients are acutely sensitive to any implication that they are feigning, the diagnosis is delivered as a real, involuntary, believed condition, never as an unmasking. Get the stance wrong and no amount of correct information will land; get it right and the explanation becomes therapeutic.

TRAP **Confronting the patient with the psychological nature of the attacks.** The classic error is to announce that the tests are normal, that there is nothing physically wrong, and that the problem is psychological, phrasing the patient hears as an accusation of pretending. It ruptures engagement at the exact moment it was most needed. The exam answer is that reassurance of genuineness comes first and the word feigning never appears.

37.3 What a good explanation contains

Two complementary templates, one shared skeleton. The content of the disclosure has been formalised into overlapping frameworks that a candidate can reproduce. Kaplan and Sadock's Comprehensive Textbook of Psychiatry organises it into a **four**-element **framework**: reassure the patient that the symptoms are genuine and not fabricated; explain how the diagnosis was reached and give the condition its names; offer the hypothetical mechanisms, including a neurobiological model, together with any identified risk factors; and review the treatments that are effective. Lishman's Organic Psychiatry sets out a closely matched non-confrontational explanation template that runs along the same axis but foregrounds the concept of **dissociation** as the model offered to the patient: explain why the events are not epilepsy and what the patient does have; reassure the patient that they are not suspected of putting the attacks on and that the disorder is very common; discuss the causes and maintaining factors; and outline the treatment. Read side by side, the two are one method described twice, and holding them together lets the reader see both the structure and the substance of a competent explanation.

BEAT OF THE EXPLANATION	COMPREHENSIVE TEXTBOOK OF PSYCHIATRY FRAMEWORK	LISHMAN'S EXPLANATION TEMPLATE
Establish genuineness	reassure the symptoms are real and not fabricated (element one)	reassure the patient is not suspected of putting on the attacks, and that the disorder is very common
Name and account for it	explain how the diagnosis was reached and give the condition its names	explain why it is not epilepsy and what the patient does have, described as dissociation
Offer a mechanism	give hypothetical mechanisms including a neurobiological model, plus identified risk factors	discuss the causes and maintaining factors
Point to the way out	review the effective treatments	outline the treatment

The two authoritative templates for delivering the diagnosis, aligned beat for beat; they are one non-confrontational method, with the Comprehensive Textbook of Psychiatry naming four discrete elements and Lishman foregrounding dissociation as the model given to the patient.

Two features of the shared skeleton are worth drawing out because they carry the therapeutic load. The first is that naming the condition is not optional politeness but an active ingredient: telling the patient the names of what they have, and that it is common, converts a frightening private experience into a recognised entity with a known course and a treatment. The second is the mechanism step. Offering a neurobiological model, a way of understanding how distress and dissociation can produce a physical event without conscious intent, gives the patient a face-saving and biologically respectable account that neither blames them nor mystifies the symptom. This is what makes the explanation non-accusatory in substance rather than only in tone.

MNEMONIC

The four R's

The Comprehensive Textbook of Psychiatry explanation in four beats. **R**eal: reassure the symptoms are genuine and not fabricated. **R**eached: explain how the diagnosis was made and give the condition its names. **R**eason: offer the hypothetical mechanisms and a neurobiological model, with any risk factors. **R**emedy: review the treatments that work. Each beat is something the framework actually states, and the order is deliberate because genuineness has to be settled before anything else is heard.

37.4 Withdrawing antiepileptic drugs safely

The drugs come off, but they come off carefully. Many of these patients arrive already taking one or more antiseizure medications, prescribed empirically over years for what was assumed to be refractory epilepsy. A logical consequence of the diagnosis is that treatment aimed at epilepsy is unnecessary and can be stopped, but this cannot be done abruptly. Patients must be cautioned that any **antiepileptic drug withdrawal** should be gradual and that the drugs should not be stopped suddenly. There are two reasons, and both should be made explicit in the consultation. First, the original diagnosis may not have been secure, and abrupt cessation risks precipitating genuine seizures or, at worst, status epilepticus if any true epileptic tendency is present or

coexists. Second, several antiseizure drugs carry their own withdrawal phenomena and are unsafe to discontinue quickly regardless of the indication. The tapering is therefore a medical decision made with the neurologist rather than a message to act on alone, and it is best framed to the patient as a planned, supervised reduction rather than as stopping useless tablets. Carers should be brought in early, with the patient's consent, both because they often witness the attacks and manage them and because they will be central to supporting the withdrawal and the psychological work that follows.

GOOD TO REMEMBER

When you tell a patient their events are not epilepsy, do not let them walk out and stop their antiseizure medication that evening. Spell out that the withdrawal will be gradual and supervised, arrange it with the neurologist, and involve a carer early with the patient's consent. The safe-taper message is part of the diagnostic conversation, not an afterthought to it.

37.5 Psychological therapy is the mainstay

The definitive treatment is a talking treatment. Once engagement is secured, the mainstay of treatment is psychological therapy, and the evidence base points most clearly to **cognitive behavioural therapy**. Kaplan and Sadock's Comprehensive Textbook of Psychiatry and Lishman's Organic Psychiatry both place structured psychotherapy at the centre of management, with randomised controlled trials, including multicentre trials, supporting cognitive behavioural therapy through improvements on measures of functioning and symptom severity. The therapy is not generic. The protocols used have been modelled on the cognitive behavioural approaches developed for panic disorder and for post-traumatic stress disorder, which is coherent with the phenomenology: the attacks often have the abrupt, fear-driven, dissociative quality of a panic or trauma response, and the therapeutic work addresses the triggers, the catastrophic appraisals and the avoidance that maintain them. The clinical message is that referral for a specific, evidence-based psychological therapy is the definitive step after diagnosis, not a vague onward disposal, and that the diagnosis should be delivered in a way that sets up and motivates that referral. Naming cognitive behavioural therapy as the treatment during the disclosure itself is one of the ways the explanation is made to point toward a way out rather than merely closing a door.

37.6 The limited and specific place of medication

Drugs treat the company the disorder keeps, not the disorder itself. The counterpart to the primacy of psychotherapy is a genuinely restricted role for pharmacology. There is no convincing evidence that pharmacological treatment reduces the frequency or severity of the attacks themselves, and so psychopharmacology should be reserved for treating comorbid psychiatric conditions rather than aimed at the seizures. Where depression, an anxiety disorder or post-traumatic stress disorder coexists, and comorbidity is common, treating it on its own merits is appropriate and may indirectly help; the psychiatric comorbidity and predisposing factors that accompany these attacks are set out in their own topic. But the drug is being given for the depression or the anxiety, not as an anticonvulsant-by-another-name for the non-epileptic attacks. This distinction is examinable precisely because it runs against the reflex to reach for a

tablet when a patient keeps having what look like fits, and because it protects the patient from an accumulating and useless polypharmacy layered on top of the antiseizure drugs that are already being withdrawn.

> **PITFALL**
>
> The instinct to control the attacks with medication, swapping one antiseizure drug for another or adding a sedative or an antipsychotic to suppress the events, is both futile and harmful. It has no evidence behind it for the attacks themselves, it reinforces the patient's belief that the problem is a drug-responsive brain disease, and it delays the psychological treatment that actually works. Prescribe only for a diagnosed comorbid condition, and be able to say why.

37.7 Stepped, multidisciplinary management

What actually has to be organised around the patient. The elements above assemble into a stepped, **multidisciplinary** model in which definitive care is shared across **two** clinical limbs that must stay coordinated. The relevant standards of care are the epilepsy guidelines of the International League Against Epilepsy and of NICE, which frame the medical management and the withdrawal of unnecessary treatment, together with the epilepsy-psychiatry reference literature, Lishman's Organic Psychiatry and Kaplan and Sadock's Comprehensive Textbook of Psychiatry, which frame the communication and the psychological treatment. Around those, the plan can be set out on the management spine.

The **LOCUS** of care is predominantly outpatient and community-based, because this is a disorder managed through a planned sequence of consultations and a course of therapy rather than through admission; the acute presentation of an attack may bring the patient to an emergency setting, but the emergency is not where the disorder is treated, and the task there is recognition and avoidance of iatrogenic harm rather than definitive care. The **focus** moves by phase: the acute target is the diagnostic conversation itself and safe handling of any attacks, the short-term target is securing engagement and beginning the supervised withdrawal of unnecessary antiseizure medication, and the mid-to-long-term target is the course of psychological therapy and the treatment of any comorbidity, with attention throughout to function, work and driving. The **modus** spans every treatment mode. The psychological mode is the definitive one and is described above. The pharmacological mode is deliberately narrow, confined to comorbidity. The medical and procedural mode belongs to the neurologist. The social mode runs through all of it: educating and enlisting carers with consent, and addressing the practical fallout for employment and driving that the diagnosis carries.

- The **neurologist** manages any comorbid neurological conditions, tapers the unnecessary antiseizure medication, and advises the patient on work and on driving.
- The **mental health clinician** delivers the definitive psychological treatment and manages comorbid psychiatric disorder.
- Both share responsibility for the consistent, non-confrontational message, since a contradiction between the two limbs is enough to undo the engagement the disclosure achieved.

IN INDIA

Several realities sharpen this in Indian practice. Stigma around the attacks is heavy, and families may read them as possession or as deliberate misbehaviour, so the non-confrontational, genuineness-first delivery matters even more, and a diagnosis heard as faking can damage marriage prospects and standing in ways that outlast the attacks. Trained cognitive behavioural therapists are scarce outside tertiary and metropolitan centres, which means the psychiatrist's own explanation carries a larger share of the therapeutic load and a longer follow-up may substitute for formal therapy where none is accessible. Because undiagnosed fits are frequently treated empirically with antiseizure drugs, often phenobarbital, many patients arrive on medication that must be withdrawn gradually and under supervision, making the safe-taper counselling a routine and important part of the consultation. Advice on driving and on manual or shift work has real economic weight and should be given concretely.

Four	**First-line**	**Comorbidity only**	**Two limbs**
ELEMENTS OF THE DIAGNOSTIC EXPLANATION	PSYCHOLOGICAL THERAPY, CHIEFLY CBT	THE LICENCE FOR MEDICATION	NEUROLOGIST AND MENTAL HEALTH CLINICIAN

The way the diagnosis is communicated is the pivotal, potentially therapeutic step: deliver it non-confrontationally and genuineness-first, taper unnecessary antiseizure drugs gradually rather than stopping them, make evidence-based psychological therapy the mainstay while reserving medication for comorbidity, and run the care as a coordinated collaboration between neurologist and mental health clinician.

38 · Coexisting PNES and epilepsy and prognosis of PNES

Two clinical problems fold into a single examination topic here. The first is what to do when a patient proves to have both epileptic and **dissociative seizures** at once, a combination that is neither rare nor academic; the second is what becomes of the patient with psychogenic non-epileptic seizures over the years that follow the diagnosis. Both are tested because both are routinely mishandled at the bedside. The coexistence is missed whenever a confirmed diagnosis of one seizure type is allowed to explain every event the patient has, and the outcome is misjudged whenever a falling seizure count is read as a recovering patient. This is the topic in which a candidate shows they can hold two diagnoses in one patient and can keep the fate of the seizures separate from the fate of the person.

Coexistence and prognosis are linked, not merely filed together. Comorbid epilepsy is itself one of the features that darkens the outlook of the dissociative seizures, so the patient who has both is, almost by definition, in the harder group. What belongs to neighbouring topics is deliberately left to them: the definition and classification of psychogenic non-epileptic seizures within the functional neurological disorder family, and the broader conversion construct developed in the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes; the clinical features that separate a dissociative from an epileptic seizure; the video-EEG and other investigations that confirm the diagnosis; the psychiatric comorbidity and predisposing factors that underlie it; and

the way the diagnosis is delivered together with its stepped psychological and pharmacological management. What is owned here is the overlap with epilepsy and the long-term outcome.

38.1 How often the two coexist, in both directions

The overlap runs both ways, and at a similar rate. The single most useful fact to fix first is that epilepsy and dissociative seizures are not mutually exclusive diagnoses between which a clinician must choose, but conditions that frequently sit in the same patient. Kaplan and Sadock's Comprehensive Textbook of Psychiatry puts electrographically and clinically confirmed **comorbid epilepsy** at **about 10 to 20 per cent** of adults with psychogenic non-epileptic seizures, and adds that the proportion is higher still in children. Strikingly, the traffic runs the other way at much the same rate: a comparable 10 to 20 per cent of adults attending epilepsy services also have dissociative seizures. The symmetry is worth memorising, because in either clinic a substantial minority of patients carries the other diagnosis as well, and the label already in the notes is no guarantee that it accounts for the events now being described.

The direction matters because it changes what the clinician is prone to miss. In the epilepsy clinic, a patient with a secure diagnosis of epilepsy who begins to have atypical or treatment-resistant events invites the assumption that the epilepsy has worsened, when some of those events may be dissociative. In the seizure-disorder clinic the mirror error is to assume that a confirmed dissociative-seizure diagnosis rules out an epileptic fit, so that a genuinely epileptic event is waved away as psychogenic. Neither denominator, on its own, tells the clinician which mechanism produced the event in front of them. That has to be worked out event by event, which is why the coexisting group is treated as its own problem rather than as a straightforward instance of either disorder.

RULE **One seizure type never excludes the other.** A documented epilepsy does not mean that every event is epileptic, and a confirmed dissociative-seizure diagnosis does not mean the patient can never have an epileptic fit. In the coexisting patient each event must be assigned to its own mechanism on its own evidence, not deduced from the diagnosis already on file.

38.2 When the epilepsy comes first

An epileptic template that the dissociative seizures later imitate. In most patients the two conditions are simply concurrent, but in a minority the epilepsy is the older diagnosis and the dissociative seizures appear later. Lishman's Organic Psychiatry records that a pre-existing history of epilepsy precedes the onset of the dissociative seizures in **fewer than one in five** patients, a genuine minority but a clinically instructive one. The sequence matters because it hints at a mechanism. A person who has lived with epilepsy holds an intimate, embodied knowledge of what a seizure looks and feels like, and that experience can furnish the template from which a dissociative seizure is unconsciously constructed. The dissociative event is not feigned; it is modelled, without the patient's awareness, on a form they already know from the inside.

Two practical consequences follow. First, the patient with prior epilepsy whose new events do not quite match the old ones deserves careful re-appraisal rather than a reflexive increase in

medication, because the new events may be dissociative copies of the familiar fit. Second, the resemblance is precisely what makes these dissociative seizures hard to separate from the patient's own epileptic ones, since they are drawn from the same original. The discrimination of dissociative from epileptic events on their clinical features is developed in its own topic and is not repeated here; the point that belongs to this one is that a shared history gives the two event types a shared appearance, which is the deep reason the coexisting patient is so hard to read.

38.3 Why the coexisting patient is the hardest to treat

Two disorders in one patient, and a target that keeps moving. Lishman's Organic Psychiatry singles out patients with coexistent epilepsy as the most difficult group of all to manage, and the difficulty is diagnostic before it is therapeutic. Before any single event can be treated it has to be attributed to the right disorder, and in this group that attribution will not stay fixed. Showing the patient and their carers a recording of a captured attack is, in Lishman's account, invaluable, because it lets everyone see the seizure rather than argue about it; but the **semiology** of a given person's events changes over time, so a recording that settled the question a year ago may no longer describe the events happening now, and the **video-EEG telemetry** may have to be repeated. The clinician is therefore not making one diagnosis but maintaining two, and re-checking which is which as the picture shifts.

Management is guideline-anchored on the epilepsy side by the ILAE and NICE epilepsy guidelines, while the dissociative seizures draw on the epilepsy-specific psychiatric-management literature and Lishman's Organic Psychiatry; the way the diagnosis is communicated and the stepped psychological and pharmacological management of the dissociative seizures are set out in their own topic and are not reopened here. The **LOCUS** of care is principally outpatient, and for this group ideally a joint neurology and psychiatry clinic; the events themselves are best characterised on an inpatient monitoring unit when clinic assessment cannot resolve them, and the acutely misread attack may first present to the emergency department. The **focus** shifts by phase: the acute priority is safety and the avoidance of treating a dissociative event as status epilepticus; the short-to-mid-term work is the event-by-event sorting of epileptic from dissociative seizures; and the long-term aim is to keep both diagnoses simultaneously in view, so that the epilepsy is not undertreated while the dissociative seizures are not chased with ever-larger doses of medication. The **modus** spans every mode of care at once: antiepileptic optimisation for the epileptic seizures, the psychological treatment that belongs to the dissociative ones, and the shared work of showing patient and carers their own recorded events.

- Events have to be sorted one at a time, because no single investigation labels the whole patient as one disorder or the other.
- Telemetry often has to be repeated rather than done once, which is costly and not always available.
- Both disorders are managed in parallel, so undertreating the epilepsy and overtreating the dissociative seizures are simultaneous risks.

PITFALL

The commonest way to go wrong with the coexisting patient is to treat every event as refractory epilepsy. Dissociative seizures mistaken for worsening fits attract escalating antiepileptic doses and added second and third agents, and when an event is read as status epilepticus they draw benzodiazepine loading and intensive-care admission. The result is drug toxicity and iatrogenic harm inflicted for events that no antiepileptic could ever have stopped. When a patient with established epilepsy stops responding to sensible dose increases, ask whether some of the events are dissociative before reaching for the next drug.

IN INDIA

The tool Lishman leans on, repeated video-EEG telemetry, is concentrated in a handful of tertiary and academic centres and is out of reach for most patients, so in practice the dual diagnosis usually has to be built from careful serial history and witness accounts rather than confirmed on a monitoring unit. Add the long distances, the cost of prolonged monitoring and the stigma that keeps dissociative seizures from being named openly, and events are commonly mislabelled and mistreated for years before the coexistence is recognised. Because outcome turns heavily on how early the diagnosis is made, this delay is not a bureaucratic inconvenience but a direct prognostic loss.

38.4 The natural history: this does not simply pass

The prognosis is soberer than the reassurance implies. Once the diagnosis is secure and the frightening possibility of uncontrolled epilepsy excluded, it is tempting to treat psychogenic non-epileptic seizures as a benign problem that will resolve with explanation. The naturalistic outcome data do not support that optimism. In the cohort reported by Reuber and Elger in 2003, the seizures were still occurring beyond three years in about **two-thirds** of patients. For the majority, in other words, the disorder is not self-limiting; left to its own course after diagnosis it runs on for years. This is the fact that justifies active treatment rather than reassurance alone, and it is the counterpoint to the relief that so often follows the exclusion of epilepsy.

The persistence also reframes the goal. If most patients still have events years on, the meaningful questions are not whether the seizures will vanish of their own accord, for most they will not, but which patients do better and whether stopping the seizures is even the outcome that matters most to the patient's life. The long persistence is the background against which every prognostic factor should be read: the factors do not carry a patient from certain cure to certain failure, they shift the odds within a disorder that, untreated, tends to endure.

38.5 The single strongest predictor: duration of illness

Catch it early, or watch it entrench. Across the outcome literature, one variable does more work than any other. As Couprie and colleagues showed in 1995, and consistent with what is seen in other somatoform disorders, the single most important predictor of a good outcome in psychogenic non-epileptic seizures is a short **duration of illness** at the time the diagnosis is made. The longer a patient has been having dissociative seizures before anyone names them

correctly, the poorer the outlook, and the effect is strong enough that it outranks most of the other factors that intuitively seem important.

The reason is best understood as consolidation. A disorder caught soon after onset has not yet reorganised the patient's life around it: work, relationships, benefits and self-concept are not yet built upon the illness. Once those structures have formed, the seizures become load-bearing, and dismantling them is far harder than it would have been early on. This is why the strongest prognostic factor is also the most actionable one for the clinician: duration is the one variable that a prompt, confident diagnosis can still influence, whereas age at onset or the presence of comorbid epilepsy cannot be changed. The practical corollary is that delay is itself harmful, not merely a postponement of care, and that the effort spent reaching a firm diagnosis early is repaid directly in outcome. It also sharpens why the coexisting patient, in whom the diagnosis is genuinely hard and therefore often late, tends to do badly: the very difficulty of the diagnosis lengthens the illness before it is finally named.

GOOD TO REMEMBER

Of everything that shapes the outcome of dissociative seizures, the one lever still in the clinician's hand is the time to diagnosis. A patient whose events are correctly named within months of onset has a materially better prospect than one carried for years under the wrong label. Treat a confident, early diagnosis as a therapeutic act in its own right, not merely as a piece of administrative labelling.

38.6 The wider prognostic profile

Duration leads, but it does not stand alone. Around the dominant effect of illness duration sits a cluster of demographic, clinical and social factors that push the outcome one way or the other. Kaplan and Sadock's Comprehensive Textbook of Psychiatry groups them into a poor-prognosis set and a good-prognosis set, and the two are worth holding side by side, because they are in part mirror images of one underlying dimension: how entrenched, and how supported, the patient is. Which psychiatric comorbidities occur, and why they predispose to dissociative seizures, is the business of a neighbouring topic; what matters here is only that their presence worsens the outlook.

CORRELATES OF A POORER OUTCOME	CORRELATES OF A BETTER OUTCOME
comorbid epilepsy	higher educational attainment
comorbid psychiatric diagnoses	a shorter duration of illness
a dramatic seizure semiology	being in employment
a younger age at onset	**acceptance of the diagnosis** and social support

The poor-prognosis and good-prognosis correlates read as near mirror images: the more entrenched and less supported the patient, the worse the course; the better educated, employed, supported and accepting, the better.

Read across the two columns and one theme keeps returning. Comorbid epilepsy and a younger age at onset mark a more biologically and developmentally embedded problem; a dramatic semiology and a longer illness mark one more woven into the patient's identity and relationships;

while education, employment, social support and an accepting stance toward the diagnosis all describe a person with resources to draw on and a reason to recover. The list is not a scoring system but a way of forming a considered expectation, so the clinician neither promises a cure the data do not support nor writes off a supported, recently diagnosed patient with every reason to do well.

MNEMONIC

Better with EASE

The good-prognosis cluster in four beats. **E**ducation, the higher the better. **A**cceptance of the diagnosis by the patient. **S**upport, meaning social support around them. **E**mployment, being in work. Sitting above all four is the single strongest factor of all, a short duration of illness at the time of diagnosis, which the mnemonic keeps in reserve as the point to state first and forget last.

38.7 When the seizures stop but the patient does not recover

Seizure outcome and life outcome come apart. The sharpest and most examinable idea in the prognosis of dissociative seizures is that the two outcomes a clinician cares about do not move together. Kaplan and Sadock's Comprehensive Textbook of Psychiatry, echoed by Lishman's Organic Psychiatry, records that participation in psychotherapy is associated with an improvement in seizure frequency in the **one to two years** after diagnosis; the seizures, in other words, can be treated and often do get better. And yet the rates of disability do not fall over the same years, and even among patients whose seizures remit, many remain out of work and dependent on state support. The **functional outcome** lags behind, and sometimes never catches up with, the seizure outcome.

This dissociation carries a hard clinical lesson. The seizures are, in one sense, only the surface of the disorder; beneath them lie the psychological and social difficulties that generated and maintained them, and those do not dissolve merely because the visible events have stopped. A treatment plan that measures success by seizure counts alone will record a victory while the patient's life stays as constrained as before. It follows that stopping the seizures, though a real and worthwhile goal, is not the same as restoring the person, and that the rehabilitation of function, meaning work, relationships and a way out of the sick role, has to be pursued as a target in its own right rather than expected to follow once the events cease.

TRAP **Reading seizure remission as recovery.** Candidates and clinicians alike equate a fall in seizure frequency with a cure, and so predict that the patient whose events have stopped will resume an ordinary life. The outcome data say otherwise: disability persists and many remain unemployed and dependent even after the seizures have remitted. State the dissociation explicitly, that the seizures can improve while function does not, and the mark, like the patient, is secured.

10 to 20% EPILEPSY COEXISTING WITH PNES IN ADULTS	**under 20%** EPILEPSY PRECEDES THE DISSOCIATIVE SEIZURES	**two-thirds** STILL FITTING BEYOND THREE YEARS	**1 to 2 y** SEIZURE GAINS AFTER PSYCHOTHERAPY

Epilepsy and dissociative seizures coexist in a substantial minority of patients in either direction, and this dual group is the hardest to treat because each event must be assigned to its own mechanism against a semiology that drifts; the disorder usually persists for years, the strongest predictor of a good outcome is a short duration of illness at diagnosis, and seizure remission must never be mistaken for functional recovery.

Antiepileptic drugs: the psychiatric interface

39 · AEDs used as mood stabilisers: valproate, carbamazepine, lamotrigine indications and evidence

Three molecules, two clinics. Valproate, carbamazepine and lamotrigine are the drugs a candidate meets first in the epilepsy clinic and again in the mood clinic, and the examiner's favourite manoeuvre is to see whether the two encounters have been kept apart. This topic takes the antiepileptic lens on purpose. It asks which of these agents is chosen for which kind of seizure, what the controlled evidence actually says about that choice, and which adverse-effect signatures they share when they are used as anticonvulsants. The **AEDs used as mood stabilisers**, with their bipolar indications and evidence - why a psychiatrist reaches for them in bipolar illness, how their plasma levels are followed, what they do to the ovary in that setting - are set out in full in the Mood disorders: pharmacotherapy notes and are not rehearsed here. The systematic, class-level account of the anticonvulsants as psychotropic agents belongs to the **systematic psychopharmacology notes**, which in turn defer back to this chapter for the seizure-selection material.

Why should a psychiatrist care how a neurologist picks an antiepileptic drug? Because almost every decision that follows rides on that choice. The agent already on the prescription decides which psychotropics can be added without a pharmacokinetic collision, it colours mood, behaviour and cognition in ways that are easy to misattribute to the epilepsy or to a primary psychiatric disorder, and in a woman of childbearing potential it forces a conversation about pregnancy long before any pregnancy is planned. Reading the logic backwards, from the drug in hand to the seizure type it implies, is one of the more reliably rewarded skills in this part of the paper, and it starts from a single rule about how the choice is made.

39.1 Selection is anchored to the seizure type

The organising principle of antiepileptic prescribing is that the drug is matched to the seizure type and syndrome, not to the patient's psychiatric label. The framework that underwrites this, the grouping of seizures into focal, generalised and unknown onset, is the ILAE seizure classification developed in this chapter's account of how seizures are classified; the selection rules simply hang off it. The practical division a candidate must carry is between **broad-spectrum** agents, which suppress several seizure types at once, and narrow-spectrum agents, which are effective against one mechanism and may be useless, or even aggravating, against another. A broad-spectrum drug is the safe default when the syndrome is uncertain or when more than one

seizure type coexists in the same patient; a narrow-spectrum drug is a precision instrument that demands a confident diagnosis before it is deployed.

SEIZURE TYPE	FIRST-LINE ANTIEPILEPTIC	WHY IT IS CHOSEN
Focal (localisation-related epilepsy)	**carbamazepine**, with **lamotrigine** as a co-first-line option	Long-established control of focal and secondarily generalised seizures; lamotrigine holds equivalent first-line standing with a different tolerability profile
Generalised tonic-clonic seizures	**valproate**	Broad-spectrum cover across the generalised seizure types makes it the drug of choice for generalised epilepsy
Absence seizures	**ethosuximide**	Highly effective for absence, but narrow in spectrum, so it is an alternative rather than a universal generalised-epilepsy drug

First-line choice reads off the seizure type: the same three principal drugs, sorted by the kind of seizure each one is best matched to.

MNEMONIC

Focal - CBZ/LTG; Generalised - VPA; Absence - ESM

Read the spine as a coverage ladder rather than a list. Carbamazepine, with lamotrigine alongside it, answers focal onset; valproate answers the generalised syndrome by covering it broadly; and ethosuximide answers pure absence. Hold the three pairings and you can reconstruct the first move in almost any seizure-selection question without reaching for anything the fragment has not already stated.

Only once the seizure type has fixed the shortlist do the second-order considerations come into play, and these are what separate two otherwise seizure-appropriate agents from one another:

- tolerability and the everyday side-effect burden the patient has to live with;
- the pharmacokinetic interaction load the drug imposes on co-prescribed psychotropics;
- its effect on mood, behaviour and cognition;
- reproductive constraints in a woman of childbearing potential.

RULE **Match the drug to the seizure type first; let the psychiatric picture fine-tune, never overturn, that choice.** The seizure type sets the shortlist of agents that will actually control the epilepsy, and only within that shortlist do the second-order factors above decide between the candidates. A drug chosen for its psychiatric convenience but mismatched to the seizure type protects nobody, because uncontrolled seizures are the primary harm the prescription exists to prevent.

39.2 Focal epilepsy: carbamazepine, with lamotrigine alongside

Focal epilepsy, the commonest pattern of adult-onset seizures, is the setting in which carbamazepine earned its reputation as a drug of first choice, and it remains a reference agent against which newer drugs are measured. Its strength is dependable suppression of focal seizures and of the secondarily generalised convulsions that arise from them. What complicates its use is not efficacy but everything that surrounds it: it is a potent modifier of hepatic metabolism, and the pharmacokinetic interactions between antiepileptics and psychotropics that this creates, through enzyme induction and inhibition, are the province of a dedicated topic in this chapter and are only named here. For a psychiatrist, the practical upshot is that a patient arriving already on carbamazepine comes with a built-in constraint on what can be safely added.

Lamotrigine sits beside it as a genuine co-first-line agent for focal onset, chosen not because it controls seizures more forcefully but because it is generally the gentler drug to live with, better tolerated across mood and cognition in many patients. That trade-off, robust seizure control from the older agent against tolerability from the newer one, is exactly what the controlled evidence set out to quantify. Lamotrigine's own mechanism, its obligatory slow titration and the rash that titration is designed to avoid are taught where lamotrigine is covered in its own right; the point that belongs here is simply its standing as a first-line choice for focal epilepsy, level with carbamazepine on the selection map.

39.3 Generalised epilepsy: valproate and the value of breadth

Where the epilepsy is generalised, valproate is the drug of choice, and the reason is written into the word broad-spectrum. A single agent that suppresses generalised tonic-clonic seizures while also covering absence and myoclonic seizures is worth a great deal in the idiopathic generalised epilepsies, where more than one seizure type frequently coexists in the same patient and where a narrow drug aimed at only one of them would leave the others unguarded. Breadth also buys tolerance of diagnostic uncertainty: when the precise syndrome is not yet settled, a broad-spectrum drug is far less likely to aggravate an unrecognised seizure type than a precision agent chosen on an incomplete diagnosis.

That advantage is why valproate has been so hard to displace despite the difficulties attached to it, and it sets up the central tension of the whole topic, because the population with idiopathic generalised epilepsy skews young and includes many women for whom valproate is the very drug that reproductive considerations argue against. The evidence and the dilemma that flow from this are taken up in the sections that follow; here the point to fix is the pharmacology, that valproate's breadth is a real clinical asset and not merely a default. A broad-spectrum agent earns its place precisely in the patient whose seizures do not fit a single tidy category.

39.4 Absence seizures and the narrow-spectrum trap

Absence seizures are the clearest illustration of what narrow spectrum means at the bedside. Ethosuximide is a first-line drug for absence and a highly effective one, but its usefulness stops at the edge of that single seizure type. It is the textbook example of a precision agent: excellent for the target it is designed for, and no help at all outside it. The clinical consequence follows directly. In a patient whose generalised epilepsy comprises absence alone, ethosuximide is an entirely

rational and often preferable choice. In a patient who has absence together with generalised tonic-clonic seizures, the same drug controls only half the problem, and the convulsive seizures it does not touch are the more dangerous half.

When absence and convulsive seizures coexist, the broad-spectrum agent that covers both is usually the sounder single-drug choice, and ethosuximide is reserved for the pure absence picture or added to a regimen that already protects against convulsions. The exam reward here is not the name of the drug but the recognition that spectrum, not potency, decides whether a given agent is adequate as sole therapy for a given patient.

PITFALL

The recurring clinical error is to treat ethosuximide as a general-purpose generalised-epilepsy drug because it is filed under generalised seizures. It is effective for absence but does not cover generalised tonic-clonic seizures, so a patient who has both, prescribed ethosuximide alone, is left unprotected against the convulsive seizures that carry the greater risk of injury. Before ethosuximide is chosen as monotherapy, the question is always whether any convulsive seizures coexist; if they do, a broad-spectrum agent, not a narrow one, is the safer spine of the regimen.

39.5 The controlled evidence for AED selection

Selection rules are only as good as the trials behind them, and the reference study for these first-line choices is the large pragmatic comparison known as **SANAD**, run in the United Kingdom to test standard against newer antiepileptic drugs in the seizure types where a clinician genuinely hesitates. In the arm addressing focal epilepsy, carbamazepine performed well on control of the epilepsy itself, showing a better time to a sustained **12-month remission**, although that particular advantage reached statistical significance only against gabapentin among the comparators. Lamotrigine, by contrast, did better on a different and arguably more clinically honest endpoint, achieving a longer time to treatment failure.

The distinction between those two endpoints is the whole lesson of the trial and worth spelling out. Time to remission asks only whether the drug stops seizures; time to treatment failure is a composite that captures both whether the drug works and whether the patient can tolerate staying on it, since a drug abandoned for side effects has failed just as surely as one that never controlled the seizures. Lamotrigine's edge on that composite, together with its being judged cost-effective and well tolerated, is precisely why it rose from a newer alternative to a co-first-line agent standing level with carbamazepine. The examinable message is that the older drug can look stronger on raw seizure control while the newer drug wins on the measure that combines efficacy with tolerability, and real prescribing lives in that gap.

39.6 Valproate in women of childbearing potential

The single most important qualification to everything said about valproate is that its status as drug of choice for generalised epilepsy does not survive intact in women of childbearing potential. Its cosmetic and its reproductive effects together create a genuine treatment dilemma: the drug that best controls the seizures is also the drug that a woman who might become pregnant has the strongest reasons to avoid. This is not a small caveat; it actively drives selection

toward alternatives across a large and clinically common group of patients, sometimes at a real cost to seizure control. The full account of valproate's teratogenicity, the formal contraception requirements it now carries and the detail of prescribing it in women is the subject of a dedicated topic in this chapter and is not pre-empted here; what belongs to selection is the dilemma itself and the way it reshapes the first-line choice.

IN INDIA

The dilemma bites harder in Indian practice, and for structural reasons rather than clinical ones. Valproate and carbamazepine are cheap, familiar and widely stocked, so they are often the pragmatic first reach for generalised epilepsy, while the newer agents that would sidestep the reproductive problem can be costlier or less reliably available outside larger centres. The counselling and contraception apparatus that is meant to accompany valproate in a woman of childbearing potential, set out in the topic on prescribing valproate in women, is unevenly implemented, and many women are never told that their antiepileptic carries reproductive implications at all. The workable stance is to raise the question of pregnancy at the point of prescribing, not years later, and to make the selection with that conversation already in view rather than as an afterthought.

39.7 The idiosyncratic reactions that shape selection

Every one of these drugs can turn on the skin. A shared adverse-effect signature runs across the aromatic anticonvulsants and it bears directly on how they are chosen and monitored. Carbamazepine, phenytoin and phenobarbital, and more commonly lamotrigine, all carry a risk of an idiosyncratic allergic rash, occurring in **about 2 to 4 per cent** of patients on the drug and appearing characteristically early in treatment rather than after long exposure. Most such rashes are benign and settle on withdrawal, but the same agents are the ones implicated when a drug eruption escalates into its dangerous forms, **Stevens-Johnson syndrome** and **Lyell's syndrome**, the latter also known as toxic epidermal necrolysis. Because the reaction is a class phenomenon, a serious rash on one aromatic agent casts a shadow of cross-reactivity over the others, which is itself a selection constraint when a substitute has to be found.

For selection and monitoring the implications are concrete. The higher rash risk attached to lamotrigine is the reason its introduction is governed by the slow-titration rule taught where that drug is covered in its own right; here the point is the class effect and its early timing. Any patient starting one of these drugs is warned to report a rash at once, is reviewed while the risk window is still open, and has the drug stopped rather than continued through an evolving eruption.

2 to 4%	12 months
ALLERGIC RASH, EARLY ON THE AROMATIC AEDS	SANAD SUSTAINED-REMISSION ENDPOINT

GOOD TO REMEMBER

Treat any new rash in the first weeks of an aromatic antiepileptic as a stop-and-review event, not a wait-and-see one. The benign rashes and the earliest stages of the dangerous ones look alike at onset, and the safe habit is to assume the worst until the drug is off and the patient reviewed. It costs little to stop and re-choose within the seizure-appropriate shortlist; it costs a great deal to continue through the rash that turns out to be the start of Stevens-Johnson syndrome.

39.8 Choosing, switching and sharing the prescription

Antiepileptic selection is a management problem, and it is best organised the way any management problem is. The standards that anchor it are the epilepsy guidelines of the International League Against Epilepsy (ILAE) and the National Institute for Health and Care Excellence (NICE), which govern first-line choice by seizure type, together with the epilepsy-specific psychiatric-management literature that addresses the mood, behavioural and reproductive considerations layered on top of the seizure diagnosis.

The **LOCUS** of this care is overwhelmingly the outpatient interface between neurology and psychiatry, where drugs are started, cross-titrated and swapped, and where the real hazard is that two prescribers do not see each other's changes; it moves inpatient when an unstable epilepsy or a complex switch is safest under observation, and to the emergency setting when a blistering rash or a serious reaction has already declared itself. The **focus** shifts by phase: in the acute phase the target is simply matching a seizure-appropriate agent to the confirmed seizure type; across the medium term it is tolerability and the emerging mood, behavioural and cognitive footprint, which are the concern of the topics on individual-AED psychiatric side effects and on AED-attributable mood change; and in the long term it is the patient whose circumstances change, above all the woman moving toward pregnancy, for whom the selection made years earlier must be revisited. The **modus** spans every lever available: the pharmacological lever of first-line choice and rational switching within the seizure-appropriate shortlist, with the interaction pharmacokinetics, lamotrigine's titration and valproate's use in women each handled by their own topics; the procedural options for drug-resistant epilepsy, whose psychiatric consequences are covered under outcomes after epilepsy surgery; and the social and educational levers of adherence support, driving and employment counselling, and honest communication between prescribers so that a change on one chart is known to the other.

TRAP **Importing the mood-stabiliser evidence to justify the antiepileptic choice.** Because these three drugs are also mood stabilisers, candidates reach for the bipolar evidence to defend a seizure-selection decision, or the reverse. The two literatures are separate: a drug's standing in bipolar illness, developed in the Mood disorders: pharmacotherapy notes, says nothing about whether it matches a given seizure type, and the antiepileptic evidence says nothing about mood prophylaxis. Keep the streams apart and cross-refer them cleanly; conflating them is the error the paired framing is designed to expose.

Antiepileptic choice follows the seizure type, not the psychiatric label: carbamazepine or lamotrigine for focal epilepsy, valproate for generalised epilepsy, ethosuximide for absence alone, with the mood-stabiliser role of these same drugs a separate lens taught elsewhere.

40 · Psychiatric side effects of individual AEDs (levetiracetam, topiramate, phenobarbital, vigabatrin and others)

Every antiepileptic drug that quietens the cortex can also unsettle the mind, and the pattern of that disturbance is rarely random. This topic is the per-agent map of psychiatric adverse effects: which individual antiepileptic drug produces which behavioural or affective syndrome, in whom, and how often. The marks are won in the vignette where a named drug is started and a named psychiatric picture follows, the previously equable adult who turns explosively short-tempered, the child transformed into an agitated whirlwind, the patient who develops a frank psychosis within weeks of a new prescription. The examiner wants the drug identified, its signature recognised, and that recognition converted into the right clinical move. The presiding fact, worth stating at the outset, is that the offending molecule is usually the single part of the whole neuropsychiatry of epilepsy that the clinician put there and can simply take away.

Two boundaries keep this topic honest. The shared attribution problem, separating a drug effect from the disease, ranking the agents by cognitive load and holding the cognitive map apart from the behavioural one, belongs to the cognitive and behavioural adverse effects of the antiepileptic drugs, taught as their own topic in this chapter; it is invoked here, not rebuilt. The class-wide suicidality caution attached to the antiepileptic drugs as a group, with its pooled-trial evidence and regulatory history, is likewise a separate topic and is named here only to be set against the agent-specific signal this fragment genuinely owns. The mood-stabiliser reading of valproate, carbamazepine and lamotrigine, their bipolar indications, mechanism of mood action and monitoring, lives in the Mood disorders: pharmacotherapy notes and is not reopened. What remains, and what is ours, is the individual drug and the particular way it disturbs thought, mood and behaviour.

40.1 Why each agent carries its own psychiatric signature

The phenotype tends to follow the pharmacology. An antiepileptic drug does not disturb behaviour at random; the flavour of the disturbance is shaped by how the molecule alters cortical and limbic signalling. Agents that push hard on inhibitory tone can blunt affect and, paradoxically, disinhibit behaviour; agents that touch glutamatergic excitation can drive irritability and aggression; agents with a broad, multi-system pharmacology tend to produce a diffuse mix of slowing and low mood. Some of these mechanisms are firmly established for the individual drugs and worth stating precisely: **vigabatrin** acts as an irreversible GABA-transaminase inhibitor, **perampanel** is a selective AMPA-receptor antagonist, and **topiramate** is, among its several actions, a weak carbonic anhydrase inhibitor. The point of the frame is not to memorise receptor targets but to make the psychiatric signature predictable, so a clinician can reason from the drug to the symptom rather than be surprised by it.

Across the class these reactions share a common grammar more useful than any single drug fact: they appear early, are frequently dose related, and are above all reversible, settling when the offending agent is reduced or withdrawn. That grammar separates a drug reaction from the disease-driven psychopathology of epilepsy and gives the clinician a lever the underlying condition rarely offers.

RULE **A new psychiatric symptom in a treated epilepsy patient is drug-attributable until the timeline says otherwise.** Antiepileptic psychiatric adverse effects are typically early, dose related and reversible, so they are testable in a way the disease is not: map the symptom against every recent start and dose change, and if it tracks the drug, the most powerful intervention is usually to the prescription rather than a tranquilliser laid on top of it.

40.2 Levetiracetam: irritability, aggression and behavioural activation

The behavioural agent every candidate must know. Of all the modern antiepileptic drugs, **levetiracetam** is the one most reliably tied to emergent behavioural trouble, a cluster informally known as levetiracetam rage that spans **aggression**, agitation, anger, anxiety, apathy, depressed mood, emotional lability, hostility and irritability. In adults treated for epilepsy, these behavioural adverse effects were reported in **13.3% against 6.2% on placebo**, roughly a doubling over background, as documented in Kaplan and Sadock's Comprehensive Textbook of Psychiatry. They characteristically arise early, often within the first month of treatment, which is exactly when the drug timeline is easiest to reconstruct. Most cases are irritability and a shortened fuse rather than dangerous behaviour, but the reaction is real enough that it drove discontinuation for behavioural reasons in 1.7% of adults, and suicide attempts, though rare, were recorded in 0.5% against none on placebo. The teaching point is one of attribution: this is a drug-attributable phenomenon, mechanistically and clinically distinct from the disease-intrinsic aggression, irritability and episodic dyscontrol of epilepsy that is taught elsewhere in this chapter, and the two demand opposite responses, one aimed at the prescription and the other at the illness.

That distinction has therapeutic teeth. When irritability or overt aggression emerges in step with starting or up-titrating levetiracetam, the rational move is to reconsider the agent, not to reach reflexively for an antipsychotic to sedate it, because the disturbance frequently settles once the drug is lowered or replaced. The recognised amplifier is a personal history of mood or behavioural disorder; the recognised mitigation is a slow introduction, the same total dose reached gradually being far better tolerated than the dose reached in a hurry.

40.3 Levetiracetam psychosis and the paediatric amplification

Rarer psychosis, and a much louder problem in children. Beyond the behavioural cluster sits a smaller but harder-edged signal. Levetiracetam-associated psychosis occurred in **0.7% of treated adults against 0.2% on placebo**, an uncommon reaction but a genuine one, and an emergent psychosis within weeks of starting the drug should raise it as a first hypothesis rather than being ascribed reflexively to the epilepsy. What matters clinically is that a drug-induced psychosis of this kind is expected to remit on withdrawal, which sets it apart from the enduring interictal psychoses of epilepsy that are taught as their own topics in this chapter.

The behavioural liability is markedly worse in children, which is the fact most often examined and most often forgotten. In paediatric epilepsy the behavioural adverse effects, spanning agitation, depression, emotional lability, hostility and hyperkinesia, were reported in **37.6% against 18.6% on placebo**, roughly triple the adult figure, and the reaction was severe enough to force discontinuation in 10.9% of children. A drug merely irritating in an adult can therefore be genuinely disabling in a child, warranting a frank warning to the family and a lower threshold for switching. The accompanying figure sets the adult and paediatric burdens side by side, and the contrast is the whole lesson.

GOOD TO REMEMBER

Prescribe the behavioural offenders defensively from the first tablet. Start low and climb slowly, name the likely reaction to the patient and one family member in advance so it is reported early rather than silently endured, and set an explicit review point in the first few weeks when the risk is highest. In children the warning is not optional: the behavioural burden is far heavier than in adults, and a forewarned family reports the change that an unwarned family mislabels as naughtiness or as the illness.

40.4 Topiramate: cognitive slowing, low mood and rare psychosis

A distinctive drag on cognition, with a mood cost behind it. Topiramate is unusual in appearing on both the cognitive and the behavioural registers at once. Its best-known effect is neuropsychological: a psychomotor and cognitive slowing together with a characteristic word-finding and memory difficulty that patients describe vividly as the sentence that stalls and the ordinary noun that will not surface. That cognitive signature is developed in full as part of the cognitive and behavioural adverse effects of the antiepileptic drugs, taught as its own topic in this chapter; what concerns us here is topiramate's specifically psychiatric liability, which is real if less famous. Alongside the slowing, topiramate can produce depression and mood problems, and, less commonly, a frank psychosis, as tabulated in Kaplan and Sadock's Comprehensive Textbook of Psychiatry. The broad, multi-target pharmacology of the drug is the usual explanation for why cognition, language and mood are all touched rather than any one domain in isolation.

The practical consequence is a low threshold of suspicion: a patient who becomes slow, inarticulate and low in mood after topiramate is introduced is usually showing the drug, not a new primary depressive illness. Because that low mood is easily attributed to the burden of epilepsy itself, the discipline is to check the drug timeline before reaching for an antidepressant, since lowering or replacing the topiramate may resolve the picture without a second agent.

40.5 Phenobarbital: depression, irritability and the paradoxically hyperactive child

The old drug with a young patient's paradox. **Phenobarbital**, one of the oldest agents still in wide use, carries a psychiatric profile that spans depression, irritability, aggression, impaired cognition and attention and sedation, so that a dulled, low, short-tempered patient on long-standing phenobarbital is a common and correctable sight. The examinable twist is age-dependent. Where the drug sedates the adult, in children it characteristically produces the opposite, a **paradoxical hyperactivity** and behavioural activation, so that the child becomes

restless, overactive and hard to manage rather than drowsy. This paradoxical reaction is a classic teaching point precisely because it inverts the expected effect and is so readily misread. The mechanism is not the point the examiner tests; the recognition is.

IN INDIA

Phenobarbital remains a common first-line antiepileptic in children across resource-limited Indian practice, kept there by its low cost, once-daily dosing and secure supply. That makes its paradoxical behavioural activation a live problem rather than a museum fact. A previously settled child who turns overactive, inattentive and irritable after starting phenobarbital is often mislabelled as naughty, as having attention-deficit hyperactivity disorder, or as expressing the epilepsy, and is disciplined or medicated for it while the true cause sits in the drug chart. Ask when the behaviour began relative to the prescription; where it tracks the phenobarbital, revisit the agent, and where an affordable alternative exists the behavioural cost is a legitimate reason to move.

40.6 Vigabatrin: precipitated psychosis and irreversible visual-field loss

Two liabilities, and one of them you cannot take back. Vigabatrin sits among the agents most likely to precipitate psychiatric decompensation, and it precipitated psychotic and affective symptoms in between 3% and 10% of patients, a reaction made considerably more likely by a significant past psychiatric history, as set out in the Companion to Psychiatric Studies. That prior history is the screening question, since it marks the patient in whom vigabatrin most readily unmasks a psychosis. The drug's second and defining liability is not psychiatric at all but limits its use so severely that no account of it is complete without it. Vigabatrin causes a progressive, bilateral **visual field constriction** that is irreversible, and it is this vision loss, more than any behavioural effect, that has pushed its use to the margins, so that it is now largely restricted to infantile spasms, particularly those of tuberous sclerosis, as described in Goodman and Gilman's Pharmacological Basis of Therapeutics.

MNEMONIC

Vigabatrin: two things you cannot reverse

The drug irreversibly blocks GABA-transaminase and it irreversibly narrows the visual fields, and it can precipitate a psychosis on top, most readily in the patient with a psychiatric past. Hold the pairing of an irreversible enzyme action and an irreversible field defect and both halves of why the agent is used sparingly fall into place.

PITFALL

The avoidable harm is to keep a patient on vigabatrin for its undoubted efficacy while never arranging or repeating visual field assessment, so that peripheral vision is quietly and permanently lost before anyone notices. The field defect is asymptomatic until it is advanced, and it does not recover once established, which is exactly why baseline and periodic perimetry, not a wait for the patient to complain, is the standard of care wherever the drug is used.

40.7 Perampanel, zonisamide, gabapentin and the wider field

The signatures continue past the headline drugs. Perampanel earns particular attention because its behavioural liability is cleanly dose dependent, which makes it both predictable and manageable. Hostility and aggression reactions rose with dose, reported in about **12% at the 8 mg daily dose and 20% at 12 mg daily, against 6% on placebo**, generally emerging within the first six weeks of treatment, and the drug carries a boxed warning for serious neuropsychiatric reactions, as documented in Kaplan and Sadock's Comprehensive Textbook of Psychiatry. That boxed warning is agent-specific and is a different regulatory object from the class-wide suicidality caution named earlier; here it flags this drug's own propensity to hostility and aggression. Because the effect climbs with the dose, the management is written into the prescription: keep the dose as low as seizure control allows. **Zonisamide** carries its own mixed neuropsychiatric profile of agitation, depression, psychosis, impaired cognition and somnolence, a reminder that the sulphonamide-related agents are not behaviourally silent. **Gabapentin**, often thought of as benign, shows a behavioural signal in children, with hostility and aggressiveness in 5.2% against 1.3% on placebo and emotional lability or behavioural problems in 6% against 1.3%, per its prescribing information, so that even the gentler agents earn a caveat in the paediatric setting.

AGENT	SIGNATURE PSYCHIATRIC ADVERSE EFFECT	PRACTICAL NOTE
Levetiracetam	irritability, aggression and emotional lability; rarer psychosis	worst in children; early onset
Topiramate	depression and mood change; rare psychosis, alongside cognitive slowing	slow titration mitigates
Phenobarbital	depression and irritability; paradoxical hyperactivity in children	the paradox inverts by age
Vigabatrin	precipitated psychosis and affective symptoms	worse with a psychiatric past; watch the visual fields
Perampanel	dose-dependent hostility and aggression	boxed warning; keep the dose low
Zonisamide	agitation, depression, psychosis, cognitive impairment, somnolence	a mixed profile
Gabapentin	hostility and behavioural problems in children	less benign than assumed in the young

The individual-agent psychiatric signatures owned here. The point of the table is discrimination: the drug most associated with rage is not the one most associated with cognitive slowing, and the one whose worst liability is not psychiatric at all sits in the middle.

40.8 The agent-specific suicidality signal and managing the reactions

Not every antiepileptic carries the same self-harm risk, and that is the point. Set against the flat, class-wide suicidality caution taught under its own topic in this chapter, there is a sharper, agent-specific finding that belongs here. A British case-control analysis found that the antiepileptic drugs whose clinical-trial depression rate exceeded 1% carried an increased self-

harm liability, and for those agents the risk of self-harm or suicidal behaviour was raised roughly **three-fold**, with no such association for the remaining drugs, as reported by Andersohn and colleagues in 2010 and summarised in Kaufman's Clinical Neurology for Psychiatrists. The high-depression-risk agents that carried the signal were:

- levetiracetam;
- tiagabine;
- topiramate;
- vigabatrin.

The discrimination is the mark: the self-harm risk is not spread evenly across the class but concentrates in the agents most prone to lower mood, which is exactly what one would predict and which the class-level caution alone would obscure.

Managing these reactions is the most rewarding work in the neuropsychiatry of epilepsy, because the cause is one the clinician can adjust. The relevant standards are the epilepsy guidelines of the International League Against Epilepsy and of NICE for the seizure disorder itself, read alongside the epilepsy-specific psychiatric-management literature for the behavioural dimension. The **LOCUS** of care is overwhelmingly the outpatient clinic, where recognition, dose review and rationalisation proceed across visits; only a severe reaction, a frank drug-induced psychosis or emergent suicidality with a safety concern, briefly pushes care toward an emergency or inpatient setting. The focus shifts by phase. The acute task is recognition and attribution, placing the symptom on the drug timeline; the short-to-mid-term task is to act on the offending agent by lowering the dose, slowing an over-rapid titration, or switching to a behaviourally kinder drug; and the long-term task is to keep the regimen lean and to monitor mood and behaviour over time, and for vigabatrin to keep the visual fields under review. The modus spans every mode: dose reduction and substitution pharmacologically, with simple restraint in the target dose for dose-dependent offenders such as perampanel; and around these the psychological and social work of forewarning patients and families, psychoeducating households that mislabel a drug effect as character or illness, and attending to the schooling, work and driving consequences of behavioural change. The choice of an antidepressant where genuine depression persists, weighed against its effect on the seizure threshold, is handled in its own dedicated topic in this chapter.

TRAP **Reading a drug-induced psychosis as the epilepsy's own psychosis.** When a patient becomes psychotic within weeks of starting levetiracetam or vigabatrin, the tempting error is to file it as an interictal psychosis of epilepsy or as forced normalisation and to add an antipsychotic, leaving the true cause in place. The tell is the timeline: a drug-induced psychosis is time-locked to the prescription and expected to remit on withdrawal, whereas the chronic psychoses of epilepsy are not. Account for the drug before you diagnose the disease.

Levetiracetam	Vigabatrin	Perampanel	Reversible
IRRITABILITY AND RAGE	PSYCHOSIS PLUS VISION LOSS	DOSE-DEPENDENT HOSTILITY	RECOGNISE AND REDUCE

Each antiepileptic drug has its own psychiatric signature: levetiracetam drives irritability and aggression, with florid behavioural activation and rarer psychosis especially in children; topiramate brings cognitive slowing with low mood; phenobarbital sedates adults but paradoxically activates children; vigabatrin precipitates psychosis and, if missed, irreversible visual-field loss; and perampanel produces dose-dependent hostility under a boxed warning. Because every one of these reactions is early, dose related and reversible, the decisive intervention is almost always to the prescription, not a sedative laid on top of it.

41 · AED-psychotropic pharmacokinetic interactions (enzyme induction and inhibition)

Two drugs, one liver, and a level that will not sit still. This topic earns marks out of proportion to its apparent dryness because it is where the neurologist's prescription and the psychiatrist's prescription silently collide. The older antiepileptic drugs are among the most powerful modifiers of hepatic drug metabolism in the whole formulary, and the psychotropics a patient with epilepsy so often also needs are cleared down exactly the pathways those antiepileptics disturb. Get the direction of the interaction wrong and a stabilised patient relapses on an unchanged dose, or a quietly accumulating drug climbs into toxicity, with nothing in the prescription itself to show what happened. The examinable core is not a list of pairings to be memorised but a single organising question asked of every co-prescription: does this antiepileptic induce the enzymes that clear its partner, or inhibit them, and therefore does the partner's level fall or rise?

Several of these same drugs are also used as mood stabilisers, a role developed in the Mood disorders: pharmacotherapy notes; here the lens is purely pharmacokinetic, concerned with what one drug does to another's concentration and not with what either does to mood. Their own psychiatric side effects, likewise, are a separate topic in this chapter. The interaction runs in both directions and reaches well beyond psychiatry, so the safe prescriber holds the mechanism rather than a memorised table of pairs.

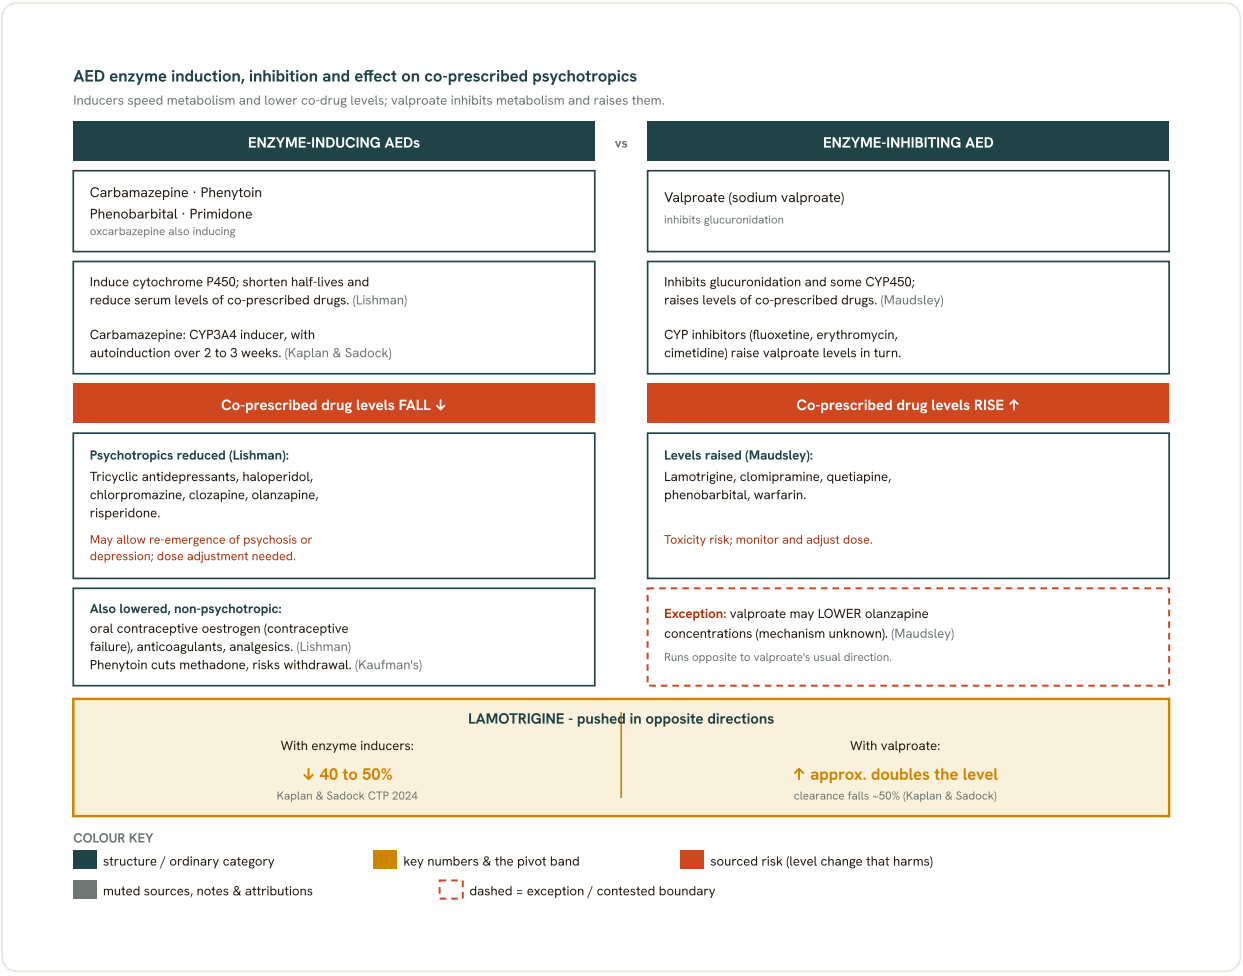


Fig 41.1 AED enzyme induction and inhibition and the effect on co-prescribed psychotropics. Enzyme-inducing AEDs (carbamazepine, phenytoin, phenobarbital, primidone) lower serum levels of many co-prescribed drugs, risking loss of efficacy, while the inhibitor valproate raises them, risking toxicity; lamotrigine is the pivot, cut 40 to 50 per cent by inducers yet roughly doubled by valproate. Attributions per box (Lishman, Kaplan & Sadock CTP 2024, Maudsley, Kaufman's).

41.1 One mechanism, two directions

Everything downstream follows from induction versus inhibition. Hepatic and gut drug metabolism runs largely through the **cytochrome P450** enzyme system, and a co-prescribed drug is only ever a substrate being processed by it. **Enzyme induction** means the antiepileptic increases the quantity of enzyme the liver manufactures, so the partner drug is metabolised faster, its half-life shortens and its steady-state serum concentration falls. **Enzyme inhibition** is the mirror image: the antiepileptic blocks the enzyme, the partner is cleared more slowly, and its level rises. The two processes even differ in tempo in a way that matters at the bedside. Induction is slow because it depends on the synthesis of new enzyme protein, so its effect builds gradually over the first weeks of co-treatment and, just as importantly, unwinds over a similar period when the inducer is withdrawn. Inhibition is comparatively rapid, appearing as soon as the inhibiting drug reaches an effective concentration. This is precisely why the whole subject collapses into one question rather than a grid to be learnt by rote: establish whether the drug in front of you is an inducer or an inhibitor, and the direction in which every one of its partners will move is already decided. The drug names are the easy part; the direction is the mark.

RULE **Inducers lower, inhibitors raise; always ask which.** Before adding, stopping or swapping any drug in a patient on antiepileptics, decide whether the antiepileptic induces or inhibits the enzymes that clear the other agent. An enzyme-inducing drug drives co-prescribed levels down and risks loss of effect; an enzyme-inhibiting drug drives them up and risks toxicity. The direction of the interaction, not the identity of the particular pair, is what you are being tested on and what keeps the patient safe.

MNEMONIC

Induce reduces, inhibit inflates

The whole interaction in four words. An enzyme-**induc**ing antiepileptic **reduc**es the level of the drug given with it; an enzyme-in**hib**iting antiepileptic in**flat**es it. Attach the four classic strong inducers, carbamazepine, phenytoin, phenobarbital and primidone, to the first half and valproate to the second, and the direction of nearly every examinable pairing falls out without a memorised table.

41.2 The enzyme-inducing antiepileptic drugs

A short, high-yield list of powerful inducers. The strong enzyme-inducing antiepileptics are almost all the older, first-generation agents, and for examination and prescribing purposes they behave as a single class with one shared consequence.

- **Carbamazepine**
- **Phenytoin**
- **Phenobarbital** and the barbiturates
- **Primidone**, which is itself metabolised to phenobarbital

Taken together, Lishman's Organic Psychiatry describes how these agents accelerate the clearance and reduce the serum levels of many co-administered drugs, an effect that spares almost nothing else the patient is taking: other antiepileptics, oral contraceptives, anticoagulants, analgesics and the psychotropics that are this chapter's concern. The clinical signature of induction is a quiet one. It is not a new symptom but the return of an old one, because the partner drug is now being cleared before it can work. A patient whose depression or psychosis had been controlled begins to slip, the antiepileptic and the psychotropic doses are both unchanged, and the cause stays invisible unless the prescriber is thinking pharmacokinetically. The breadth of that list matters as much as any single pairing, which is why the safe habit is to treat the addition of an enzyme-inducing antiepileptic as a standing reason to review every other drug on the chart.

41.3 Carbamazepine and autoinduction

The inducer that induces its own metabolism. Carbamazepine deserves a paragraph of its own because it does something the others do not do so dramatically: it induces the very enzyme that clears itself. Kaplan and Sadock's Comprehensive Textbook of Psychiatry identifies it as a **CYP3A4** inducer that undergoes **autoinduction**, a process developing over the first **2 to 3 weeks** of treatment. The practical consequence is counter-intuitive and repeatedly tested. A dose that produced a good level in the first week yields a lower level a fortnight later, because the liver

has caught up and is now metabolising the drug faster, so slightly higher doses are often needed once autoinduction is complete. There is a second lesson folded into this. Because carbamazepine's own enzyme induction takes those first weeks to reach full strength, every interaction it has with other drugs is time-dependent in the same way: its power to suppress a partner's concentration is not fully expressed on the first day but grows over the same window. A level checked or a clinical judgement made too early therefore reads falsely reassuring, and the true effect on both carbamazepine and its partners is only seen once the enzyme population has finished expanding.

GOOD TO REMEMBER

When you start carbamazepine, do not judge the steady state from the first days. The level you measure early will drift down as autoinduction proceeds, and so will the levels of anything else the patient is taking that carbamazepine induces. Recheck the drug level and the patient's mental state once the induction has settled, not while it is still building, and warn the patient that a dose which felt right at first may need a modest increase.

41.4 What induction does to psychotropic levels

The stabilised patient who relapses on an unchanged dose. The interaction that bears most directly on a psychiatrist is the fall in plasma antipsychotic and antidepressant concentrations that enzyme-inducing antiepileptics produce. Lishman's Organic Psychiatry documents the effect across both drug classes and across generations of antipsychotic, from the older agents such as haloperidol and chlorpromazine to the second-generation drugs, and across the tricyclic antidepressants. The consequence the source draws out is the clinically important one: as the concentrations drop, a previously controlled psychosis or depression can re-emerge, and the psychotropic dose may need to be raised to compensate. This is a silent relapse mechanism, because nothing about the presentation announces that a level has fallen; the illness simply returns while the prescription looks unchanged and adherence looks perfect. The correct response is not to escalate treatment blindly but to recognise the pharmacokinetic cause, and either raise the psychotropic dose against a target level or, where the choice is still open, prefer an antiepileptic that does not induce. Therapeutic drug monitoring, where the drug can be assayed, converts an invisible fall into a measured one.

41.5 Lamotrigine, the drug pulled both ways

One partner, two opposite interactions: the clearest illustration of the whole principle. No drug demonstrates the bidirectionality of antiepileptic interactions better than lamotrigine, because it is a substrate for both an inducing and an inhibiting antiepileptic and moves in opposite directions with each. Enzyme inducers drive it down. Kaplan and Sadock's Comprehensive Textbook of Psychiatry records that co-administration with inducers such as carbamazepine, phenytoin, phenobarbital or primidone lowers the lamotrigine plasma concentration by **40 to 50 per cent**. Valproate does the reverse, and does it powerfully. As an inhibitor of the glucuronidation pathway that clears lamotrigine, valproate **approximately doubles** the lamotrigine concentration. Set those two facts side by side and the same nominal

lamotrigine dose can represent twice or half the effective exposure depending purely on the company it keeps. The stakes of reading that direction correctly are high, because lamotrigine's dosing is dominated by the need to titrate slowly to avoid a serious rash, and that slow-titration rule and its rash risk are taught where lamotrigine is covered in its own right in this chapter. The pharmacokinetic fact to carry from here is its simplest summary: valproate roughly doubles the level while inducers roughly halve it, so the direction of the adjustment is set entirely by which partner the lamotrigine shares the chart with.

41.6 Valproate: the inhibitor on the antiepileptic bench

The antiepileptic that raises levels instead of lowering them. Valproate is the important counter-example to the inducing antiepileptics, and holding it firmly as an inhibitor is the second half of the whole topic. Rather than inducing cytochrome P450, it inhibits **glucuronidation**, so the drugs paired with it are cleared more slowly and their concentrations climb. The Maudsley Prescribing Guidelines note that valproate raises the levels of a spread of co-prescribed drugs: lamotrigine most strikingly, but also tricyclic antidepressants such as clomipramine, quetiapine, phenobarbital and warfarin. The last two make the point that this is no psychiatric curiosity, since a rising phenobarbital level brings sedation and a rising warfarin effect brings a bleeding risk. The traffic also runs the other way. Other enzyme inhibitors, and the same source names erythromycin, fluoxetine and cimetidine, in turn raise valproate's own concentration, so valproate is both a perpetrator and a victim of inhibition. And then the exception that tests real understanding: contrary to its usual level-raising direction, valproate significantly lowers the plasma concentration of olanzapine, by a mechanism that remains unknown. That single reversal is the detail most likely to catch a candidate who has learnt valproate as a drug that only ever pushes concentrations up.

TRAP **Assuming valproate raises every level.** Because valproate is the textbook enzyme inhibitor and famously doubles lamotrigine, candidates generalise it into a drug that can only push concentrations upward. It cannot. It significantly lowers the plasma concentration of **olanzapine**, by an undetermined mechanism, and that exception is exactly the kind of item an examiner uses to separate rote recall from a grasp of the pharmacology. The direction of an interaction is a property of the specific pair, not a fixed label stamped on one drug.

41.7 Interactions beyond the psychotropics

The induction that defeats a pregnancy plan, an anticoagulant or an opioid substitution. The enzyme-inducing antiepileptics do not confine their reach to psychiatric drugs, and three non-psychiatric interactions are high-yield because each carries a discrete and serious harm. Lishman's Organic Psychiatry notes that induction reduces the efficacy of oral contraceptives, oral anticoagulants and analgesics through the same acceleration of metabolism. The first of these is the reproductive one. Kaplan and Sadock's Comprehensive Textbook of Psychiatry warns that carbamazepine, and enzyme-inducing antiepileptics generally including oxcarbazepine, lower estrogen levels and can cause contraceptive failure, so higher-dose estrogen preparations or a reliable non-hormonal method are needed. The distinct and much larger problem of valproate's teratogenicity and the formal contraception requirements it imposes on women of childbearing

potential belong to the topic on prescribing valproate in women covered elsewhere in this chapter; the point here is only the pharmacokinetic one, that induction can silently disarm an oral contraceptive.

The second is opioid substitution. Kaufman's Clinical Neurology for Psychiatrists describes how phenytoin's induction reduces the activity of **methadone** and can precipitate opioid withdrawal, which is why the methadone dose should be raised in anticipation before phenytoin is started rather than after the patient becomes unwell. The third, anticoagulation, follows the same arithmetic: a patient stabilised on warfarin can drift to a subtherapeutic anticoagulant effect when an inducing antiepileptic is added, and back toward a bleeding risk if that antiepileptic is later withdrawn and the induction unwinds.

IN INDIA

These pharmacokinetics are not academic in Indian practice, because the cheap and widely available first-line antiepileptics are precisely the strong inducers: carbamazepine, phenytoin and, in resource-limited settings where it remains a genuine first-line drug, phenobarbital. A large share of patients with epilepsy are therefore on an enzyme inducer, so the induction interactions dominate the ground: reduced psychotropic efficacy in a population with high psychiatric comorbidity, and **contraceptive failure** in women who may never have been told that their antiepileptic can defeat the pill. Where the budget allows, a non-inducing antiepileptic sidesteps the whole problem at once. Where it does not, the workable answer is active review of co-medication and explicit contraceptive counselling, rather than assuming the two prescriptions are independent of each other.

PITFALL

The interaction does not only appear when an inducing antiepileptic is started; it reverses, with equal clinical force, when one is stopped. Reduce or discontinue an enzyme-inducing antiepileptic in a patient who is also on a psychotropic, an anticoagulant or a contraceptive, and the co-drug that had been cleared too fast is now cleared normally again, so its level climbs over the following weeks toward, and sometimes past, the toxic range. Clinicians who remember to adjust doses when adding an inducer routinely forget to readjust them when withdrawing one. Every induction interaction is a two-way street, and the return journey is the one that is missed.

ANTIEPILEPTIC AND EFFECT	PARTNER DRUG LEVEL	REPRESENTATIVE PARTNERS AND THE RESULTING RISK
Inducers on psychotropics (induction)	falls	antipsychotics and tricyclic antidepressants become subtherapeutic; silent relapse
Inducers on other drugs (induction)	efficacy falls	oral contraceptives, oral anticoagulants and analgesics lose effect
Carbamazepine on contraception (induction)	estrogen falls	contraceptive failure; use higher-dose estrogen or a non-hormonal method
Phenytoin on methadone (induction)	activity falls	opioid withdrawal; raise methadone before starting phenytoin
Valproate on most partners (inhibition)	rises	lamotrigine, quetiapine, warfarin, phenobarbital; toxicity and bleeding
Valproate on olanzapine (exception)	falls	the one direction reversal to remember

A direction map of the examinable interactions: read the middle column first, because the direction of the level, not the identity of the pair, is what the answer turns on.

41.8 Prescribing safely around the interactions

The mechanism is only useful if it changes the prescription. There is no single "treatment" of a pharmacokinetic interaction; the management is a disciplined prescribing method, and it is best organised the way any management problem is. The standards that anchor it are the epilepsy guidelines of the International League Against Epilepsy (ILAE) and the National Institute for Health and Care Excellence (NICE), which govern antiepileptic selection and the preference for non-inducing agents where a co-prescription makes that sensible, together with the Maudsley Prescribing Guidelines for the psychotropic side of the same chart.

The **LOCUS** of this care is overwhelmingly the outpatient clinic, where medications are added, stopped and cross-titrated, and where the real danger is that the two prescribers, often a neurologist and a psychiatrist, do not see each other's changes. It moves to the inpatient setting when a complex cross-titration is safest under observation, and to the emergency department when an interaction has already declared itself as frank toxicity or as withdrawal. The **focus** shifts by phase. At the moment of any change, the target is anticipation: predict the direction of the interaction and adjust pre-emptively rather than reactively. Across the first weeks, the target is the slow-building nature of induction, so a level or a mental state assessed too early gives false reassurance and must be repeated once the effect is fully expressed. In the longer term, the target is vigilance at every subsequent change to either drug, because each addition or withdrawal reopens the interaction afresh. The **modus** spans all the levers available: the pharmacological lever of dose adjustment in the predicted direction, or the cleaner option of selecting a non-inducing antiepileptic so the interaction never arises; the procedural lever of therapeutic drug

monitoring, which turns an invisible interaction into a measured quantity for the drugs that can be assayed; and the social and educational levers of explicit counselling, about contraception, adherence and never stopping a drug abruptly, and above all communication between prescribers so that a change on one chart is known to the other. Managed this way, the interaction becomes predictable arithmetic rather than an ambush.

2 to 3 weeks	40 to 50%	about double
CARBAMAZEPINE AUTOINDUCTION DEVELOPS	FALL IN LAMOTRIGINE WITH INDUCERS	LAMOTRIGINE LEVEL WITH VALPROATE

Enzyme-inducing antiepileptics (carbamazepine, phenytoin, phenobarbital, primidone) lower the levels of the drugs given with them, while valproate, an inhibitor, raises them, with olanzapine the notable exception, so the direction of the interaction, not the drug's name, must drive every dose decision.

42 · Lamotrigine: mechanism, titration, rash and bipolar maintenance role

The antiepileptic whose dose has to be earned slowly. Lamotrigine is the antiepileptic a psychiatrist meets more often than almost any other, because it sits at the crossing point of two clinics: it is a first-line drug for focal epilepsy and, separately, a familiar name in the mood clinic. This topic takes the antiepileptic lens deliberately. It asks how the drug works against seizures, why its introduction is governed by an unusually strict titration schedule, and why that schedule exists at all, which is to say the rash that has shaped every prescribing rule attached to the molecule. The maintenance targets below are the representative destinations of that slow titration, not a substitute for the age-specific and product-specific escalation steps of the current label. The maintenance mood-stabiliser role in bipolar disorder, for which lamotrigine is equally well known, is set out in full in the Mood disorders: pharmacotherapy notes and is only named here; none of that mood-stabiliser evidence is rehearsed in this account.

The marks on this topic are almost entirely practical and almost entirely about safety. An examiner is rarely interested in whether a candidate can recite that lamotrigine blocks a sodium channel; the interest is in whether the candidate can start the drug without harming the patient. That means knowing how the starting dose and the speed of escalation change with whatever else the patient is already taking, and knowing the one adverse reaction that turns a well-tolerated drug into a dangerous one. Everything below is organised around that single clinical spine.

42.1 Antiepileptic mechanism of action

A membrane-stabilising sodium channel blocker that restrains glutamate release. Lamotrigine's anticonvulsant action is that of a **voltage-gated sodium channel blocker**. Kaplan and Sadock's Comprehensive Textbook of Psychiatry describes it binding preferentially to the channel in its inactivated state, so that neurons which are firing rapidly and repetitively, the pattern of a seizure, are the ones whose membranes are most effectively stabilised, while normal background firing is comparatively spared. The downstream consequence, and the part most

often asked, is presynaptic: by damping this pathological high-frequency firing the drug reduces the release of the principal excitatory neurotransmitter, **glutamate**. That combination, a use-dependent sodium channel action expressed as a fall in excitatory transmission, is the mechanistic sentence to carry. It is worth holding apart from the quite different set of ideas invoked to explain why the same molecule stabilises mood, which belong to the notes named above; here the concern is purely the antiepileptic one, and the sodium channel is where it begins.

42.2 The serious rash and its boxed warning

The reaction that defines the drug. The reason lamotrigine cannot simply be started at a therapeutic dose is a risk of serious, potentially life-threatening skin reaction grave enough to carry a **boxed warning**. In premarketing trials the incidence of serious rash was about **1 per cent** in children under sixteen years and about 0.3 per cent in adults, the higher paediatric figure being one reason for particular caution at the younger end. The reactions that give the warning its weight are the severe cutaneous drug eruptions, **Stevens-Johnson syndrome** and, at its most extreme, **toxic epidermal necrolysis**, and rash-related deaths have been reported. These are not common events, but they are catastrophic when they occur, and their possibility is what converts an otherwise well-tolerated drug into one whose introduction demands a specific discipline. The examinable point is less the exact figure than the category of risk: a drug that can, rarely, kill through the skin, and that therefore has to be started in a way that keeps that risk as low as it can be made.

42.3 Benign rash and the stop-at-first-sign rule

Common, usually harmless, but never to be watched. Most rashes on lamotrigine are not the dangerous kind. A benign rash occurs in **up to 10 per cent** of patients, typically a simple maculopapular eruption appearing early in treatment, and the great majority of these settle uneventfully once the drug is withdrawn. The difficulty is that at their onset the trivial rash and the first hours of a severe one can look alike, and there is no reliable way at the bedside to be certain which is beginning. Kaplan and Sadock's Comprehensive Textbook of Psychiatry gives the resulting rule its characteristically cautious form: the drug should ordinarily be discontinued at the first sign of any rash, unless the rash is clearly not drug-related. That default toward stopping is deliberately conservative. It accepts that many benign rashes will lead to a discontinuation that was, in hindsight, not strictly necessary, because the cost of that caution is small and the cost of continuing through the one rash that turns out to be Stevens-Johnson syndrome is not. In practice the patient is told before ever taking the first tablet to report any rash at once, rather than wait to see whether it worsens.

GOOD TO REMEMBER

Before you write the first lamotrigine dose, find out what else the patient is already taking, because that single question decides the whole starting regimen. A patient on valproate needs a much lower, slower start; a patient on an enzyme-inducing antiepileptic needs a higher one, taken up more quickly. Choosing the starting dose without knowing the co-medication is the commonest way to get a lamotrigine prescription wrong before the patient has swallowed anything.

42.4 The timing window of the dangerous rash

When the risk is live, and when it has largely passed. The serious rash is not a hazard spread evenly across the life of the prescription; it is concentrated early. Nearly all life-threatening rashes appear within the first weeks of treatment, the great majority falling **between 2 and 8 weeks** after the drug is started, with only isolated cases reported later. This early clustering is what makes the titration period the critical window, and it explains why review is front-loaded into the first weeks rather than spread evenly across the year. It also gives the reassurance that matters for the long-term patient: a person who has climbed slowly through the escalation schedule and reached a stable maintenance dose without a rash has passed through the interval of greatest danger, and the risk on continued steady dosing is very much lower. The corollary, developed below, is that anything which effectively returns the patient to the beginning of treatment, most importantly a break in dosing, reopens that early window and has to be treated as a fresh start rather than a simple continuation.

42.5 Why the rash happens: risk factors and the slow-titration rule

The rash is a rate phenomenon, not a dose phenomenon. The single idea that unifies lamotrigine prescribing is that the risk of a serious rash tracks the speed at which the serum concentration is allowed to rise far more than the dose eventually reached. A high maintenance dose arrived at gently is safer than a modest dose reached in a rush. Three circumstances raise the risk, and each is a way of making the level climb too fast:

- escalating the dose more rapidly than the recommended schedule allows
- starting above the recommended initial dose
- **concurrent valproate**, which pushes the lamotrigine level up and so amplifies any given rate of increase

From this single mechanism the entire prescribing discipline follows. Because the danger is the rate of rise, the protection is a slow rise, and Kaplan and Sadock's Comprehensive Textbook of Psychiatry together with Lishman's Organic Psychiatry make **gradual, stepwise introduction** the defining rule for starting the drug. It is the reason lamotrigine takes weeks rather than days to reach a useful dose, an inconvenience that is not a shortcoming of the schedule but its entire purpose.

RULE **Introduce lamotrigine slowly; the rate of rise, not the final dose, sets the rash risk.** A gradual dose introduction is the crucial manoeuvre for attenuating the risk of a serious rash, and it is the one prescribing rule that must never be compromised for the sake of reaching a therapeutic effect sooner. Every other adjustment, for co-medication or for a missed-dose restart, is a variation on protecting that slow rise.

MNEMONIC

Low, slow, and stop at the first rash

The whole safety of the drug in five words. Start **low**, climb **slow**, and **stop** at the first sign of any rash rather than watching it. The first two protect against the rash by keeping the rate of rise gentle; the third protects against it by refusing to gamble on a given rash being the benign kind. Nothing in the sentence has to be memorised as a number, which is exactly why it survives the pressure of a real clinic.

42.6 Titration and target doses by co-medication

The starting dose is a function of the company the drug keeps. Because the rash risk is set by the rate at which the level rises, and because other drugs powerfully change that level, lamotrigine has no single correct starting dose or target; both depend on the co-medication. Lamotrigine is cleared chiefly by hepatic **glucuronidation**, and it is this pathway that its partners disturb, so its elimination half-life shifts markedly with what else is on the chart. The detailed pharmacokinetics of enzyme induction and inhibition are the subject of the topic on antiepileptic and psychotropic interactions and are not reworked here; the point that belongs to lamotrigine itself is the practical consequence for how it is started and where it is aimed.

With valproate the level roughly doubles, so the standard response, drawn from Lishman's Organic Psychiatry and Kaplan and Sadock's Comprehensive Textbook of Psychiatry, is to halve the usual start and climb more gently still. With an enzyme-inducing antiepileptic the opposite holds, and the drug can be both begun higher and taken up more briskly. The accompanying table sets the three situations side by side.

CO-MEDICATION CONTEXT	STARTING APPROACH AND TITRATION	TARGET MAINTENANCE DOSE	APPROXIMATE HALF-LIFE
Lamotrigine monotherapy	standard gradual introduction	200 mg/day	about 28 hours
With valproate	halved start, 25 mg on alternate days, slower climb	100 mg/day	about 56 hours (roughly doubled)
With enzyme-inducing antiepileptics	start 50 mg/day and titrate up more quickly	400 mg/day	about 14 hours (roughly halved)

Read the table across a single row: the co-medication decides the start, the target and the half-life together, which is why the first clinical question is always what else the patient takes.

PITFALL

The commonest serious prescribing error with lamotrigine is to start it at the ordinary dose in a patient already taking valproate. Valproate roughly doubles the lamotrigine concentration, so a standard start becomes, in effect, a rapid escalation, driving the level up exactly as fast as the rash risk fears. The valproate combination is one of the named risk factors for a serious rash for precisely this reason, and it demands a halved, slower start rather than the usual one. The mirror error, under-dosing the patient on an enzyme inducer and then wrongly concluding the drug does not work, is far less dangerous but wastes months of treatment.

42.7 Restarting after an interruption

A break in treatment is a return to the beginning. One of the most clinically important and most often forgotten rules attaches to the patient who has stopped the drug for a while, whether through non-adherence, a supply problem, or an admission in which doses were missed. The safe accommodation that the slow titration built up is lost when the drug is absent, so restarting at the previous maintenance dose reproduces exactly the rapid rise in level that the original slow titration was designed to avoid. Kaplan and Sadock's Comprehensive Textbook of Psychiatry gives the operational rule: if the drug has not been taken for **5 days or more**, it should be reintroduced using the full, gradual initial escalation regimen from the start, not resumed at the dose the patient was previously on. Serious rashes have followed a rapid restart, which is why this rule is not a bureaucratic formality but a genuine safety instruction.

TRAP **Resuming at the old maintenance dose after a gap.** The instinct, when a patient who was stable on lamotrigine returns after missing it, is to restart them on the dose they were taking. If the interruption has lasted beyond that short threshold, that instinct is dangerous: the drug must be re-titrated from the beginning of the escalation schedule, because the accommodation built by the original slow climb is gone and a full-dose restart reopens the early window of serious-rash risk. Examiners favour this item precisely because the reflex answer is the wrong one.

IN INDIA

Lamotrigine is a more expensive antiepileptic than the older first-line agents that dominate Indian practice, and its obligatory slow titration adds a second, less obvious cost: it takes weeks of graded dosing, and therefore several clinic contacts, before the drug reaches a useful level. In a system where follow-up can be irregular and paid out of pocket, this creates real temptations to shortcut the schedule, to restart at the old dose after a lapse rather than begin again, or to abandon a drug that seems to be doing nothing during its long run-in. Each of those shortcuts is the specific behaviour the rash rules exist to prevent. Where lamotrigine is genuinely indicated, the counselling that goes with it, that the slow start is the safety and not a delay, and that a rash means stop and call, is as much a part of the prescription as the tablets themselves.

42.8 Prescribing lamotrigine safely

The mechanism only helps if it changes the prescription. There is no separate treatment of the lamotrigine rash beyond preventing it and responding decisively when it appears, so the management here is a disciplined prescribing method. The standards that anchor it are the

epilepsy guidelines of the International League Against Epilepsy (ILAE) and the National Institute for Health and Care Excellence (NICE), which govern antiepileptic selection and safe use, together with the drug's own boxed-warning prescribing information, which is the primary authority on the rash and its titration schedule.

The **LOCUS** of care is overwhelmingly the outpatient clinic, where the drug is started and titrated over weeks and where the whole safety of the exercise is decided; it moves to the inpatient setting only when an unstable epilepsy or a complex cross-titration is safest under observation, and to the emergency and dermatology setting the moment a severe rash or a systemic reaction declares itself, at which point the drug is stopped and the patient managed as a dermatological emergency. The **focus** shifts by phase. In the acute phase, the first weeks, the target is simply to climb the escalation schedule slowly and to keep review close, because this is the interval in which the dangerous rash lives. Across the medium term the target is reaching an effective maintenance dose appropriate to the co-medication without ever sacrificing the slow rate of rise. In the long term the targets are adherence and the correct handling of any interruption, since a break resets the patient to the beginning. The **modus** spans every available lever: the pharmacological lever of a starting dose and titration matched to the co-medication, halved and slowed for valproate and raised for enzyme inducers; the procedural lever of immediate discontinuation and supportive dermatological care should a serious rash appear; and the social and educational levers of counselling the patient to report any rash at once, never to restart at the old dose after a lapse, and to understand that the slow introduction is the very thing keeping them safe.

about 1%	up to 10%	2 to 8 weeks	5 days+
SERIOUS RASH, CHILDREN UNDER 16 IN TRIALS	BENIGN RASH ON LAMOTRIGINE	WINDOW FOR THE LIFE-THREATENING RASH	GAP THAT FORCES RE-TITRATION FROM THE START

Lamotrigine is all about the rash: start low and titrate slowly, halve the starting dose with valproate and raise it with enzyme inducers, stop at the first sign of any rash, and re-titrate from the beginning after a break of 5 days or more, because the rate at which the serum level rises, not the final dose, decides the risk of a life-threatening reaction.

43 · Valproate teratogenicity, contraception requirements and prescribing in women

No prescribing decision in the epilepsy clinic is more heavily freighted than whether to place, or to leave, a woman of childbearing potential on sodium valproate. For a given patient valproate may be the most effective drug she has for her seizures, and yet it is the most teratogenic antiepileptic in routine use, harming a developing fetus in ways both structural and behavioural. That combination, high efficacy against seizures set against high risk to the fetus, is precisely what makes the topic a perennial examination favourite and a live regulatory concern. The woman with epilepsy cannot be managed as though the drug could simply be withdrawn the moment a pregnancy test turns positive: uncontrolled seizures carry their own maternal and fetal

dangers, so the clinician is forced into a genuine risk-against-risk calculus rather than a simple instruction to stop.

The examinable core has four parts, and a strong answer holds them apart: the size and shape of the teratogenic signal, its dependence on dose, the neurodevelopmental teratogenicity that no anomaly scan can exclude, and the framework of contraception, preconception optimisation and antenatal surveillance that now governs prescribing. The mood-stabiliser and polycystic-ovary aspects of valproate belong to the Mood disorders: pharmacotherapy notes and are only named here; what follows is the epilepsy-facing account.

The most economical way to see valproate's position is to compare drugs on a single broad endpoint. Taking **any adverse outcome**, meaning either a major congenital malformation or fetal death, the exposed-pregnancy rates separate the agents unambiguously, and valproate stands far above the alternatives:

20.3%	10.7%	8.2%	1.0%
VALPROATE, ANY ADVERSE OUTCOME	PHENYTOIN	CARBAMAZEPINE	LAMOTRIGINE

43.1 The teratogenic signal and where valproate sits among the antiepileptics

A class risk with a wide internal gradient, and valproate at the wrong end of it. Antiepileptic drugs as a group raise the baseline risk of major congenital malformation. Pooled across the older agents, Lishman's Organic Psychiatry puts the malformation rate in exposed pregnancies at about **5 to 9 per cent**, roughly twice the background-population rate, so a woman on any antiepileptic begins pregnancy with an elevated absolute risk before the specific molecule is chosen. The clinically decisive point, and the one an examiner is testing, is that this class figure conceals a broad gradient between drugs, and valproate occupies the unfavourable extreme.

Prospective pregnancy registers have made that gradient concrete. In **the UK Epilepsy and Pregnancy Register**, drawn on by Lishman's Organic Psychiatry, the major-malformation rate on valproate monotherapy was substantially higher than on the common alternatives, and the ordering is the pattern to carry into the exam:

MONOTHERAPY	MAJOR MALFORMATION RATE
Valproate	6.2 per cent
Lamotrigine	3.2 per cent
Carbamazepine	2.2 per cent

On the same register, denominator and endpoint, valproate roughly doubles the risk seen with lamotrigine and nearly trebles that with carbamazepine.

Two features matter beyond the digits. The ordering is consistent across datasets and endpoints: whether the outcome is a narrow count of major malformations or the broader composite shown at the head of this account, valproate is the outlier and the alternatives sit well below it. And register figures describe monotherapy under controlled conditions, so real-world risk is pushed higher still by two further variables, dose and combination, taken in turn below. The lamotrigine

and carbamazepine numbers appear only as comparators that fix valproate's position; their own prescribing stories are told elsewhere in this chapter.

RULE **Valproate carries the highest teratogenic risk of the drugs in routine antiepileptic use, and that risk is dose-dependent.** Every downstream decision, whether to start it at all, how to counsel, what dose to accept and how intensively to screen, follows from these two facts held together. Miss either one and the management falls apart.

43.2 Structural malformations: neural tube defects and the wider pattern

The signature lesion is a defect of neural-tube closure. The malformation most specifically associated with intrauterine valproate is the **neural tube defect**, classically **spina bifida**, and the risk in valproate-exposed pregnancies is of the order of about **1 to 2 per cent**. That figure must be read against the natural history of the defect. The neural tube closes very early in the first trimester, ordinarily before a woman knows she is pregnant, which is the single most important reason that prevention before conception matters far more than anything done once a pregnancy is confirmed. By the time the drug's risk is discussed in the antenatal clinic, the window in which the defect is determined has usually closed.

The anomaly profile is broader than the neural tube alone. Kaplan and Sadock's Comprehensive Textbook of Psychiatry records that valproate exposure also predisposes to cardiac and to skeletal malformations, alongside a recognisable pattern of dysmorphic features, so the structural picture spans several systems rather than one:

- neural tube defects, above all spina bifida, the lesion most tightly linked to valproate
- cardiac malformations
- skeletal and limb anomalies, together with the characteristic facial dysmorphology of the exposed infant

Why valproate does this is only partly understood. The molecule interferes with folate metabolism, inhibits histone deacetylase and generates oxidative stress, each of which can disturb the tightly timed processes of neural-tube closure, cardiac septation and neuronal migration. None is established as the sole cause, but together they explain the drug's reach across several developmental systems and why folate supplementation, addressed under management, has a rational place even though it cannot abolish the risk.

43.3 Dose-dependence and the special danger of polytherapy

Teratogenicity that scales with exposure, not an all-or-nothing property of the molecule. The harm valproate does is not a fixed toll levied on every exposed pregnancy equally; it rises with the amount of drug the fetus sees. Risk climbs with the daily maternal dose, and the inflection examiners look for is a threshold in the region above **1000 mg per day**, beyond which the malformation rate increases steeply. This dose-dependence is among the most reliably examined facts in the area, because it converts a blunt warning into a graded instruction: if valproate cannot be avoided, the objective is the lowest dose that still controls the seizures, since a lower dose is genuinely less dangerous even though no dose is safe.

Combination therapy compounds the problem. A regimen containing valproate alongside other antiepileptics carries a malformation rate near 9 per cent, higher than valproate monotherapy and than most other polytherapies, so a valproate-containing combination is the pattern to be most reluctant to create or to leave in place in a woman who might conceive. The two levers point the same way: keep the dose as low as seizure control allows and the drug alone rather than in company, since simplifying away from a valproate combination is among the most valuable single changes that can be made before a pregnancy.

43.4 Behavioural teratogenicity and the fetal valproate syndrome

The harm that a normal scan cannot exclude. The most important idea in the topic, and the one candidates most often miss, is that valproate damages the developing brain in ways that have nothing to do with any visible malformation. This **behavioural teratogenicity** is the core of what is termed the **fetal valproate syndrome**, and it is present in children who were structurally normal at birth and who passed every anomaly scan. Kaplan and Sadock's Comprehensive Textbook of Psychiatry sets out a consistent and worrying picture across three overlapping outcomes.

The first is a loss of general cognitive ability. Prenatal valproate is associated with a mean fall in childhood IQ of about **10 points**, and exposed children scored roughly 9 points below lamotrigine-exposed children when assessed at three years of age, with some falling into the range of intellectual disability. The second is confirmed in population data: a nationwide **Danish birth cohort** showed roughly a 4.5-fold increase in intellectual disability, rising to about a 6-fold increase when delayed developmental milestones were counted alongside it. The third is a raised risk of **autism spectrum disorder**, with a many-fold higher risk of childhood or lifetime autism reported and supported by animal models. Taken together these define a neurodevelopmental toll that unfolds over years:

- lowered global cognitive ability, sometimes reaching the range of intellectual disability
- frank intellectual disability and delayed developmental milestones
- autism spectrum disorder and autistic traits

The clinical force here is a single asymmetry. A structural defect such as spina bifida can, in principle, be detected before birth; a lowered IQ, an intellectual disability or an emerging autism cannot. Behavioural teratogenicity is not screenable and declares itself only as the child grows, so it is the part of valproate's harm that surveillance can never reach and that only prevention can prevent.

TRAP **Treating valproate teratogenicity as a structural problem that a reassuring anomaly scan can rule out.** The neurodevelopmental harm, reduced IQ, intellectual disability and autism, occurs in children with no visible malformation and cannot be found antenatally, so a normal scan does not mean a normal outcome. Candidates who equate teratogenicity with spina bifida forfeit the point the examiner is actually testing.

43.5 Why the woman with epilepsy cannot simply stop the drug

The dilemma that puts this topic in an epilepsy chapter at all. A woman prescribed valproate for a mood indication can often be moved to a different agent without much loss, but the woman with epilepsy is frequently the one for whom valproate is the drug that actually stops her seizures, and stopping it is not a neutral act. A generalised tonic-clonic seizure in pregnancy risks maternal injury, hypoxia and status epilepticus, each threatening the fetus in its own right, so the clinician weighs a quantified fetal risk from the drug against a real, if less easily counted, risk from uncontrolled seizures, and neither side of the balance can be waved away.

There is a further subtlety drawn out by Kaplan and Sadock's Comprehensive Textbook of Psychiatry: the birth-defect and developmental-delay risks appear more prominent in **women with epilepsy** taking valproate than in those given it for bipolar disorder. The interpretation calls for caution, since the epilepsy population differs in dose, in the frequency of polytherapy, in the disorder itself and possibly in genetic liability, so this is not a clean drug-versus-drug contrast. The practical message is unambiguous: the population in which valproate is most useful is also the one in which its risks are most evident, which is why the decision belongs before conception, taken deliberately with the woman, not in the crisis of an unplanned pregnancy.

> **PITFALL**
>
> Abruptly stopping valproate the instant a pregnancy is confirmed feels protective and is often the reverse. By the time most women know they are pregnant the neural tube has already closed, so the impulsive withdrawal cannot undo the exposure that has already happened, while the sudden loss of seizure control can precipitate breakthrough seizures or status epilepticus that endanger both mother and fetus. The correct response to an unplanned pregnancy on valproate is an urgent specialist review, not a unilateral stop.

43.6 Contraception and the Pregnancy Prevention Programme

Because the greatest harm precedes recognition, contraception becomes a condition of prescribing. The logic of everything above is that the damage is largely done before a pregnancy is known, so the only reliable protection is to prevent an unintended pregnancy while the drug is taken. The regulatory response, set out in the Maudsley Prescribing Guidelines, is formal: valproate should not be used in a girl or woman of childbearing potential unless a **Pregnancy Prevention Programme** is in place, effective contraception is mandatory throughout treatment, pregnancy is to be actively avoided while the drug is taken, and some authorities go further and advise against valproate in women under fifty altogether. This reframes the prescription: valproate in a woman who could conceive is not a routine start but a decision that has to be justified and that comes bundled with a contraceptive plan.

Which method is used is a clinical judgement favouring reliable, ideally long-acting contraception, and the interaction between enzyme-inducing antiepileptics and hormonal contraception is a real but separate consideration belonging to the pharmacokinetics of antiepileptic interactions rather than to valproate teratogenicity as such. What stays with this

topic is that the contraceptive conversation is not a one-off at the first prescription but a standing part of treatment for as long as the drug and the possibility of pregnancy coexist.

MNEMONIC

No Programme, no prescription

The single sentence that governs valproate in women of childbearing potential. The drug is not started, and ideally not continued, in a woman who could become pregnant unless a formal Pregnancy Prevention Programme with effective contraception is in place. Nothing in the rule is a number to be memorised, which is exactly why it survives the pressure of a real clinic.

GOOD TO REMEMBER

At every review of a woman of childbearing potential who is taking valproate, re-ask about contraception and about any plan or wish for pregnancy, and document that you did. Teratogenic risk is a conversation to be renewed, not a consent form signed once and filed, because contraception fails and intentions change, and the one review at which the question is skipped is the one before the unplanned conception.

IN INDIA

Valproate remains widely and cheaply prescribed to women in Indian practice, often where the formal apparatus of a Pregnancy Prevention Programme is not implemented and where the drug was started by a non-specialist before any discussion of pregnancy. Many women present already pregnant on valproate, sometimes soon after marriage, at a point where the exposure has already occurred. The workable response is opportunistic rather than regulatory: raise contraception and preconception planning at every contact with a woman of childbearing potential on the drug, prefer a less teratogenic agent wherever seizure control allows, and ensure access to folate for those who continue it, since the ideal of a pre-planned switch is frequently overtaken by events.

43.7 Prescribing when valproate cannot be avoided: preconception optimisation and antenatal surveillance

When a woman's seizures respond only to valproate, the task changes from avoidance to damage limitation, best carried out as a deliberate, staged plan. The standards that frame it are the Medicines and Healthcare products Regulatory Agency (MHRA) valproate pregnancy-prevention requirements, which set the contraceptive and consent conditions; the epilepsy guidance of the National Institute for Health and Care Excellence (NICE); and the recommendations of the International League Against Epilepsy (ILAE) on women of childbearing age. Together they turn the evidence above into a routine.

The **LOCUS** of care is shared and mostly ambulatory: the outpatient neurology or epilepsy clinic, alongside the psychiatric or primary-care setting where the drug may first have been started, with preconception counselling ideally completed before any pregnancy and specialist obstetric care taking over once a woman conceives; the emergency setting enters only when a woman presents already pregnant or with breakthrough seizures, and its rule is urgent specialist review rather than abrupt withdrawal. The **focus** shifts by phase: before conception, switch to a less teratogenic

agent where possible or else reduce valproate to the lowest effective dose and simplify away from polytherapy, with folate begun in advance; at the discovery of pregnancy, continue surveillance without a dangerous abrupt stop; across the antenatal months, structural screening; and in the long term, developmental follow-up of the child, because behavioural teratogenicity emerges only with time. The **modus** spans every mode. The pharmacological mode is the lowest effective valproate dose, avoidance of valproate-containing combinations, and periconceptional folic acid at **at least 1 mg** per day with vitamins B6 and B12. The psychological mode is a documented, unhurried discussion of the risks. The procedural mode is the antenatal surveillance set out below. The social mode is secure access to contraception, to folate and to planned rather than crisis-driven care.

Surveillance, where exposure has occurred or continues, rests on **fetal anomaly ultrasound**, maternal serum screening and, where indicated, amniotic fluid sampling, timed across the pregnancy as follows:

INVESTIGATION	TIMING	PURPOSE
Fetal anomaly ultrasound	18 to 20 weeks	detect neural tube and other structural defects
Maternal serum alpha-fetoprotein	mid-trimester	raised in the presence of an open neural tube defect
Amniocentesis (amniotic fluid sampling)	16 weeks, if indicated	confirm a suspected open defect

The surveillance addresses structural defects only: useful for the neural tube but, as above, unable to detect or predict the behavioural teratogenicity that is the drug's least reversible harm.

The whole plan is front-loaded toward the decision made before pregnancy. Folate and ultrasound are worth doing, but they mitigate a risk best managed by never incurring it, which is why the contraceptive framework and the preconception switch carry more weight than any antenatal test.

> Valproate is the most teratogenic antiepileptic in common use, its harm is dose-dependent and both structural and neurodevelopmental, and it must not be given to a woman of childbearing potential unless a Pregnancy Prevention Programme with effective contraception is in place; where unavoidable, use the lowest effective dose alone, add periconceptional folate, and remember that the behavioural harm cannot be screened for at all.

Consolidate and apply

44 · Worked vignettes

A worked case is where this chapter stops being a catalogue and becomes a clinical skill. Every syndrome in the preceding sections was defined in isolation, but no patient arrives pre-

labelled. What the examiner hands you, on paper and far more so in the viva, is a short clinical picture and a demand that you place it, attribute it, and exclude what it is not before you commit to a name. These vignettes rehearse exactly that traffic between a presentation and a diagnosis. They deliberately withhold the numbers, intervals and eponyms that belong to the individual topics, because the point of a worked case is not to re-state a fact you have read but to practise the reasoning that retrieves it at the right moment.

Read each vignette as a closed problem first. Resist the pull to answer the surface, and instead run the disciplined sequence the chapter installed at the outset: place the symptom in time, weigh whether the disease, its drug or a non-epileptic process is its author, and name the imitator you have excluded. The cases below are chosen because each one baits a fast, wrong, surface answer and rewards a slower, structured one, which is precisely the discrimination that separates a pass from a distinction.

- Each turns on a single resolving datum, a timing or a drug change or a recording, rather than on the overall impression.
- Each pairs a syndrome with a look-alike the examiner expects you to exclude by name.
- Each withholds the owned specific, so that what is practised is the retrieval and not the recitation.

44.1 The man who was lucid, then was not

A cluster of convulsions, a clear interval, then a florid disturbance. A young man with long-standing epilepsy is brought to the emergency department by his family after a run of tonic-clonic seizures more frequent than usual. He recovered his orientation between and after the fits. Then, after a symptom-free interval in which he was awake, coherent and correctly oriented, he became floridly psychotic: grandiose and religiose, agitated, with paranoid conviction and disturbed sleep. On examination now there is none of the clouding of consciousness that follows a fit directly, and the family fear he has lost his mind for good.

The disciplined reading places the disturbance on the **peri-ictal axis** before it reaches for a psychiatric label. The presence of a lucid, oriented interval between the seizure cluster and the onset of psychosis is the single most valuable datum in the history, because it is what separates a **postictal psychosis** from the confusion and delirium that come on immediately after a fit with no clear-headed gap. This is not a new, unrelated psychotic illness, and it is not the chronic interictal psychosis that develops on a settled seizure baseline; it is a peri-ictal event, characteristically self-limiting, that follows a flurry of seizures across a lucid gap. The length of that interval, the usual duration of the episode and its management are developed where postictal psychosis is taught and should be recalled from there.

GOOD TO REMEMBER

A psychosis that erupts after a lucid gap following a seizure cluster is, until shown otherwise, a postictal phenomenon, and the immediate task is not a lifelong antipsychotic commitment but safety: the patient can be grandiose, agitated and at real risk to himself during a self-limiting episode. Secure him, treat the acute disturbance, coordinate with the seizure team, and let the course declare itself before you reclassify a peri-ictal event as a standing psychotic illness.

44.2 The seizures stopped and the trouble started

Good control on paper, a disturbed patient in the room. A patient whose epilepsy had been difficult for years is finally brought under tight control after a change of regimen. The seizure diary is clean, the treating team is pleased, and the recordings that were once busy with discharges have quietened. Precisely as this happens, the patient becomes psychotic, or in other presentations irritable, dysphoric and behaviourally disturbed. A clinician who has read the seizures as the whole problem is baffled: the fits are better than they have ever been, so why is the patient worse?

The counter-intuitive answer this chapter has already named is **forced normalisation**: the emergence of psychiatric disturbance, classically a psychosis, as seizure activity is suppressed and the record normalises. Its clinical cousin is the notion of an alternative psychosis, in which fits and psychosis seem to alternate rather than to coincide. The examinable move is to recognise that improving seizure control is itself on the differential for the new disturbance, so that the correct response is to rethink the regimen with the neurologist rather than to treat the psychosis as an unrelated event and drive the antiseizure drugs harder. The historical eponym attached to this phenomenon, and the mechanism proposed for it, belong to the section that develops it.

44.3 Two kinds of fear that are not the same disorder

The same emotion, one a seizure and one an anxiety disorder. Two patients describe fear. The first has brief, stereotyped, out-of-the-blue surges of intense dread that arrive without a trigger, last a short and remarkably constant time, and are sometimes accompanied by a rising epigastric sensation and a momentary lapse of full awareness with automatisms, after which she is briefly muddled. The second has longer episodes of fear that build in feared situations, come with catastrophic thoughts about dying or losing control, and leave no clouding of awareness behind them.

The temptation is to call both panic and be done. The disciplined reading uses timing and form to separate an **ictal fear**, which is the seizure itself arising from limbic structures, from an interictal panic phenomenon, which is an anxiety disorder in the settled periods between seizures. The ictal event is brief, stereotyped, unprovoked and frequently stamped with the other features of a focal seizure, the epigastric aura, the depersonalisation, the automatisms and the postictal muddle, whereas the interictal attack is longer, situationally shaped, cognitively elaborated and followed by clear consciousness. Getting this apart matters because one is treated as epilepsy and one as an anxiety disorder, and the ictal amygdala phenomenon itself is developed where ictal fear and the ictal affective states are taught.

44.4 The events that monitoring could not catch as seizures

"Drug-resistant epilepsy" that resists because it is not epilepsy. A young woman is referred with frequent, dramatic episodes labelled drug-resistant seizures despite escalating antiseizure medication. The events are prolonged, wax and wane, involve side-to-side head movements and asynchronous limb thrashing, her eyes are firmly closed and resist opening, and afterwards she recalls parts of the episode and is not deeply confused. There is a heavy load of psychological adversity in the background. Prolonged recording captures a typical event with no accompanying electrical seizure.

This is the great imitator of the epilepsy clinic, a **psychogenic non-epileptic seizure**, a functional event that mimics a fit and is confirmed not by any single bedside sign but by weighing the semiology, the recording and the psychiatric context together. The features that point away from an epileptic seizure, and the confirmatory role and limits of video-electroencephalography and the timed serum prolactin, are developed where the diagnosis is taught, as is the stepped way the news is delivered and the event is managed. As a general functional neurological disorder, it belongs within the family described in the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes. The vignette's lesson is narrower and vital: escalating antiseizure drugs against events that are not seizures does harm, and the possibility of a functional event belongs on the differential of every apparently refractory case, not only after every drug has failed.

IN INDIA

Functional seizure disorders in Indian practice frequently present late and heavily stigmatised, framed by the family in spiritual or moral terms and taken first to faith healers rather than to a clinic, so the diagnosis lands on a patient already convinced her body has failed her. Definitive video-electroencephalography sits in tertiary centres alone, so much of the work of separating a functional event from a true seizure is done from a careful history, and the manner in which the diagnosis is explained carries even more weight than in the settings where the evidence base was written.

44.5 Low mood, and two different stories about suicide

A common, treatable, under-recognised depression, and a risk that arrives from two directions. A patient with established epilepsy is low, anhedonic and irritable, sleeping poorly, with a mood picture that does not keep tidily to the shape of a classical depressive episode and that fluctuates in a way the family put down, reasonably enough, to the strain of living with fits. He is not being screened.

Two examinable points collide in this case. The first is that depression is among the commonest and most missed psychiatric comorbidities of epilepsy, that its phenomenology can depart from the textbook episode, and that it warrants active screening with an instrument designed for the epilepsy setting. The second is **attribution** applied to suicide risk, which in epilepsy runs along two separate tracks that must never be merged: a risk intrinsic to the disease, elevated on its own account against the general population, and, quite separately, a class-level signal that the antiseizure drugs may raise suicidality, flagged by a drug regulator and demanding its own

vigilance and counselling. The magnitude of the intrinsic risk, the identity of the epilepsy-specific screen, and the regulator behind the class warning belong to the sections that own them.

PITFALL

The habitual error here is to read low mood in a patient with epilepsy as an understandable reaction to a hard life and to stop there, screening no further. Depression in epilepsy is common, treatable and frequently atypical, and the suicide risk sits above that of the general population before any drug is even considered, so "of course he is low, look at what he lives with" is precisely the sentence that lets a treatable and dangerous condition go unaddressed. Normalising the mood is not the same as assessing it.

44.6 Is it the epilepsy, or is it the tablet?

The question that recurs more often than any other in this field is whose fault a symptom is. A patient becomes irritable, low and cognitively slowed some time after a change to the antiseizure regimen. The family attribute it confidently to the epilepsy, to more seizures, or to the brain disease itself, and ask for reassurance. The cross-sectional picture cannot settle it: apathy, slowed thinking and mood change look the same whether the disease or the drug produced them.

Working the case means holding both authors open and interrogating the one the family have not suspected. The temporal link to a recent drug change or dose increase, the known behavioural and cognitive profile of the particular agent, and the response to adjusting it are what break the tie, because several antiseizure drugs carry recognised psychiatric and cognitive adverse effects that are dose-related and reversible. The individual offending agents and their signature effects, and the cognitive burden attributable to the epilepsy and to repeated seizures themselves, are deliberately taught as two separate topics, so that the reader must assign the picture to disease or to drug rather than to a single comfortable cause. A woman of childbearing potential complicates it further, because the agent most implicated in mood questions may also be the one whose use in pregnancy is most fraught; the mood-stabiliser lens on these drugs, their bipolar indications and monitoring, belongs to the Mood disorders: pharmacotherapy notes, while this chapter keeps only the antiseizure lens.

44.7 The diary, the deity and the sticky conversation

A personality, changed slowly, that the examiner wants named. A patient with long-standing temporal lobe epilepsy is described by his family as profoundly changed over years rather than in any single episode. He writes compulsively and at length, filling notebooks with detailed private reflections; his life has acquired an intense religious and philosophical preoccupation; his conversation has become circumstantial and hard to bring to a close, so that visitors cannot easily leave; and his sexual interest has altered. He is not psychotic and not acutely disturbed; this is a durable interictal change of the person.

The vignette is a portrait of the eponymous interictal behavioural cluster associated with temporal lobe epilepsy, the excessive writing, the deepened religiosity, the stickiness of thought and social engagement, and the altered sexuality, which a minority of such patients develop and which the examiner expects you to recognise from the description and name. It sits against the

broader, commoner **interictal personality change** of which **viscosity**, that same stickiness of thought and difficulty disengaging, is the marker. The two are held apart as a general phenomenon and a named syndrome, and the anatomy of the temporal lobe that underlies both is developed where temporal lobe epilepsy is taught. Naming the cluster is the mark here; reconstructing the eponym from memory, rather than recalling it from the section that owns it, is where candidates stumble.

44.8 Working the next one: turning cases into a habit

The unseen vignette is the same problem in new clothes. The cases above were chosen to cover the peri-ictal syndromes, the psychoses, mood and suicide, the functional imitator and the drugs, but their value is not as a handful of facts to memorise. It is that they share a single method, and once that method is a reflex the unseen case becomes tractable, because it is only ever a fresh surface over one of a small set of discriminations. The accompanying table indexes the worked cases by the surface answer each baits and the disciplined move that resolves it; read it as a rehearsal, not as a summary of the syndromes themselves, which are owned and developed by their own sections.

THE PICTURE PRESENTED	THE FAST, SURFACE ANSWER	THE DISCIPLINED MOVE THAT RESOLVES IT
Florid psychosis after a seizure cluster	"a new psychotic illness"	find the lucid interval and place it as a postictal event
Disturbance as the fits come under control	"an unrelated psychosis, push the drugs"	consider forced normalisation and rethink the regimen
Sudden surges of intense fear	"panic disorder"	separate an ictal fear from an interictal anxiety by timing and form
Dramatic, drug-resistant "seizures"	"escalate the antiseizure drugs"	weigh a functional event and confirm it with monitoring
Low mood and thoughts of self-harm	"an understandable reaction"	screen actively and separate disease-intrinsic from drug-attributed risk
Irritability and slowing after a drug change	"the epilepsy is worsening"	interrogate the drug before blaming the disease
A slowly, deeply changed personality	"he has simply aged"	recognise the interictal cluster and name it

An index of the worked cases: the resolving move is the examinable skill; the specifics it retrieves are owned and developed by the individual topics and are not restated here.

Two errors recur across all of these, and a single mnemonic then gathers the discriminations the cases turned on.

RULE **Reason the case; do not recognise it.** A vignette rewards the clinician who runs a sequence, timing then attribution then exclusion, over the one who pattern-matches the surface to the nearest remembered syndrome. Pattern recognition is fast and, in this field, reliably wrong, because the examiner builds the case so that its surface points one way and its resolving datum points another. The mark is in the datum, not the impression.

TRAP **Answering the phenomenology instead of the diagnosis.** Candidates name what the picture looks like, a psychosis, a panic attack, a depression, a seizure, and stop, when the examiner has deliberately made the surface appearance and the true diagnosis diverge. The fear that is a seizure, the psychosis that is a postictal event, the seizure that is functional: each is missed by the reader who answers what it resembles rather than what, once placed and attributed, it actually is.

MNEMONIC

Every case is a PAIR to prise apart

The vignette payoff in one word: each classic question hides two look-alikes that must be separated. **P**ostictal states, the delirium that follows a fit directly against the psychosis that follows across a lucid gap. **A**ttribution, the disease against its drug, in mood, cognition and suicidality alike. **I**ctal versus interictal, the fear that is a seizure against the anxiety that is not. **R**eal versus functional, the epileptic seizure against its psychogenic imitator. Prise the pair apart and the case is answered; every letter names only a discrimination this chapter has already drawn.

Do not answer what the vignette looks like; find the one datum, the timing to the seizure, the drug chart, the recording, that tells you what it is, and name that.

Lucid gap	Drug chart	Last fit	The recording
POSTICTAL PSYCHOSIS, NOT DELIRIUM	THE DRUG, NOT THE DISEASE	WHERE THE SYMPTOM SITS IN TIME	A SEIZURE, OR A FUNCTIONAL EVENT

45 · Consolidation and quick review

This closing section teaches nothing new; it makes what came before retrievable. A chapter on epilepsy and psychiatry is unusually easy to read and unusually hard to recall, because its facts arrive as a long series of syndromes that look alike on the surface and separate only on detail. The examiner exploits exactly that. Questions rarely ask you to define a single entity in isolation; they hand you a patient and ask you to place the disturbance, exclude the imitator, and decide who is to blame, the disease or the drug. What follows is built for that skill. It gathers the chapter into one reasoning frame, then hands you an answer-free set of prompts to test against, so revision becomes active retrieval, not a comfortable re-reading that flatters the memory.

The prompts here deliberately withhold their answers. Each names a construct developed earlier and asks you to supply the number, interval, criterion or eponym that belongs to it; a prompt you cannot answer is a gap to close before the examination, and far better found here than in the hall. Use this section last, once the individual topics are read.

45.1 How to use this review: test, do not re-read

The method matters as much as the material. The strongest predictor of what survives to the examination is not how many times a passage has been read but how many times it has been retrieved from a blank page. **Active recall**, the effortful pulling of an answer from memory before checking it, builds durable access in a way that re-reading, which feels productive precisely because it is easy, does not. Pair it with **interleaving**: rather than drilling all the psychoses together and then all the mood syndromes, shuffle the prompts so a postictal question is followed by a valproate question is followed by a personality one, because the examination itself is interleaved and the skill tested is discrimination across categories, not recall within one.

Spaced, repeated passes compound the gain: a syndrome retrieved today, again after a few days and again after a week, is far more secure than one crammed the night before. The prompts below are short and answer-free.

- Cover the rest of the page and attempt the prompt aloud or on paper before you look anything up.
- For every number you cannot produce, mark it and return to the owning topic, not to this list.
- Shuffle the domains between passes so that no answer is primed by the one before it.
- Re-test the items you missed after a short delay, not immediately, since immediate success is mostly short-term memory.

RULE **Retrieve before you review.** Reading a fact again and recognising it is not the same as being able to produce it unprompted, and the examination only rewards production. Treat every prompt in this section as a closed-book attempt first; the checking comes second. A fact you had to fight for is a fact you will keep.

45.2 The one frame that ties the chapter together

Almost every question in this field is a variation on three questions asked in order. Faced with a psychiatric presentation in a patient known to have epilepsy, the disciplined clinician does not begin with the psychiatric differential at all. The first question is *when*: where does the disturbance sit in relation to the seizure? Placing it on the **peri-ictal axis**, the temporal framework introduced at the front of this chapter and owned there in full, narrows a sprawling differential to a short candidate list, because each position on that axis carries only a handful of syndromes. The second is *whose*: the epilepsy itself, an **attribution** to the antiseizure drug, or an independent psychiatric disorder that happens to co-occur? That one decision separates the cognitive slowing of the disease from that of the medication, and the mood change of the illness from the class-effect signal of the drug, and a large share of marks turns on it. The third is *what else*: what imitator must be excluded before the label is accepted, whether a **psychogenic non-epileptic seizure** masquerading as a fit or a primary psychiatric disorder masquerading as an interictal syndrome?

Held together, these three questions are the spine of the chapter, and the individual topics are their elaborations: a postictal psychosis is a *when* answer, the class warning on the antiseizure drugs is a *whose* answer, and the separation of a seizure from a functional event is a *what else*

answer. Learned as a list of syndromes the chapter feels endless; learned as three questions with named answers it becomes examinable.

MNEMONIC

SEIZE the timeline

The whole diagnostic approach, held as one word. **S**ituate the symptom in time, its relation to the seizure, on the peri-ictal axis. **E**xclude the imitator: a functional event for a fit, a primary disorder for an interictal state. **I**dentify the phenomenology: fear, psychosis, dysphoria, personality change or cognitive decline. **Z**ero in on attribution: the epilepsy, the drug, or an independent disorder. **E**nsure safety and treat: gauge the suicide risk and choose agents with the seizure threshold in mind. Each letter names only a step this chapter has already taught.

In a patient with epilepsy, no psychiatric symptom is fully described until you have fixed its timing to the seizure and decided whether the disease, the drug, or an independent disorder produced it.

45.3 The peri-ictal axis: what you must be able to recall

The spine of the chapter, tested position by position. The peri-ictal band is where the most easily confused material lives, because the syndromes clustered around the seizure share a phenomenology and separate mainly on their timing, their duration, and the presence or absence of a symptom-free gap after the fit. The prompts walk the axis from before the seizure to the long interictal baseline; supply the withheld specifics yourself, since each belongs to a topic developed earlier and is deliberately not repeated here.

- State the temporal windows of the peri-ictal axis in order, and place each named syndrome on it without looking.
- Give the duration over which preictal dysphoria and irritability typically build before a seizure.
- For postictal psychosis, recall the length of the characteristic lucid interval that precedes it, how long the episode usually lasts, and the seizure pattern that most often precedes it.
- Distinguish postictal confusion and delirium from postictal psychosis on the single feature of the intervening lucid period.
- Recall the phenomenology of ictal fear and ictal panic and the one feature that separates them from an interictal panic attack.
- Name the interictal mood syndrome of epilepsy characterised by an intermittent, pleomorphic dysphoria, and list its component symptoms.

TRAP **Treating the psychoses of epilepsy as one entity.** Candidates lose marks by answering a question about ictal, postictal or chronic interictal psychosis as though the three were interchangeable, when what the examiner is testing is precisely their separation. They are distinguished by their timing to the seizure, by whether a lucid interval intervenes, by their duration, and by the state of the record between attacks. Fix which psychosis is being asked about before you write a word about its features.

45.4 The psychoses of epilepsy and forced normalisation

Three psychoses, plus the paradox. Beyond the peri-ictal band sits the chronic, schizophrenia-like psychosis that arises in the interictal baseline, and alongside it **forced normalisation**, the counter-intuitive phenomenon in which behavioural disturbance emerges as seizures are suppressed and the record quietens. The examinable skill is holding these apart from one another and from a primary psychotic illness, and the prompts target exactly the discriminators you are expected to produce.

- Recall the clinical features of the schizophrenia-like psychosis of epilepsy and the seizure history against which it typically develops.
- Recall the historical eponym for forced normalisation and how it relates to the notion of an alternative psychosis.
- Give the three features that mark a psychosis as epileptic rather than as a primary psychotic disorder.
- State the antipsychotic considerations specific to a patient with epilepsy, and where the seizure-threshold ranking of agents is developed.

The last prompt points across a boundary examiners like to probe. The discrimination of epileptic psychosis from primary schizophrenia is owned by its own topic, and the phenomenology of the primary illness belongs to a separate chapter, the Schizophrenia: aetiology, phenomenology, classification and course notes, which is where the comparator should be read rather than reconstructed from memory. This chapter owns the epilepsy-related psychoses and their separation from a primary disorder; it does not own the primary disorder itself.

GOOD TO REMEMBER

When a patient's fits have just come under good control and a psychosis then appears, do not reflexively read it as a new, unrelated illness. The improving seizure control may itself be driving the disturbance, and the presentation belongs on the differential of the alternative-psychosis phenomenon before it is called anything else. A psychosis that arrives as the seizures recede is a prompt to rethink the regimen, not to press on with it unchanged.

45.5 Mood, anxiety and suicide: the disease and the drug

The domain where attribution decides the answer. Depression is among the commonest psychiatric comorbidities of epilepsy and among the most under-recognised, and its phenomenology does not always keep to the textbook shape of a primary depressive episode.

Suicide risk is elevated in epilepsy on its own account, and, separately, a regulatory signal implicates the antiseizure drugs as a class. These two suicidality stories must never be conflated: one is intrinsic to the disease, the other attributed to the medication. The prompts keep them apart.

- Recall the screening instrument designed specifically to detect depression in epilepsy and what makes it suited to that setting.
- Give the magnitude of the elevated suicide risk in epilepsy expressed against the general population, and the factors that raise it further.
- Name the regulator that issued the class warning on antiseizure-drug-attributable suicidality, and how that signal changes prescribing and monitoring.
- State how the phenomenology of depression in epilepsy can depart from a classical major depressive episode.
- Recall the anxiety presentations of epilepsy and the one that is most often mistaken for a primary panic disorder.

Antidepressant choice here follows the same seizure-threshold logic as antipsychotic choice, and the mechanism behind that logic is owned by its own topic rather than repeated across every prescribing decision. When a prompt asks which agents are safest, the answer is an application of that shared principle, not a separate list to memorise in isolation.

45.6 Personality, behaviour and cognition

Where disease and drug are hardest to tell apart. The interictal personality changes of temporal lobe epilepsy, the eponymous behavioural cluster a minority develop, the episodic aggression that is far more often interictal and non-directed than a true ictal act, and the cognitive trajectory of the illness all sit in this domain, and all force the same question: is this the epilepsy, or is it the medication? That attribution is the whole examination point, and it is why the cognitive effects of the disease and the cognitive adverse effects of the drugs are taught as two separate topics that must not be collapsed into one.

- List the components of the eponymous interictal behavioural syndrome of temporal lobe epilepsy, and name it.
- Distinguish the general interictal personality change, of which viscosity is the marker, from that named syndrome.
- Recall the pattern of aggression and episodic dyscontrol in epilepsy and why a directed, purposeful assault is rarely a true ictal phenomenon.
- Separate the cognitive effects attributable to the epilepsy and to repeated seizures from the cognitive and behavioural adverse effects of the antiseizure drugs.
- Recall the psychiatric outcomes, both improvement and the de novo disorders, that can follow epilepsy surgery.

PITFALL

The commonest clinical error in this domain is single-cause thinking: attributing a patient's mental slowing, irritability or low mood entirely to the epilepsy while the antiseizure drug quietly drives it, or the reverse. Before the disease is blamed, the drug and its dose, its recent changes and its known behavioural profile must be interrogated, and before the drug is blamed, the seizure burden and the underlying lesion must be weighed. Attribution here is rarely clean; the discipline is holding both candidates open until one is excluded.

45.7 Functional events, and the chapters next door

The great imitator, and the material this chapter borrows. A psychogenic non-epileptic seizure is the imitator the epilepsy clinic must exclude, and the skill is not a single sign but a weighing of the semiology, the investigations and the psychiatric context. The prompts below cover that ground and then point outward: several specifics a candidate reaches for in this field are owned by other chapters, and it is worth knowing where they live so the borrowed material is read at its source rather than half-remembered here.

- Give the clinical features that most reliably separate a functional non-epileptic event from an epileptic seizure at the bedside.
- Recall the role of **video-electroencephalography** and of the timed serum prolactin, and the limits of each, in confirming the diagnosis.
- State how the diagnosis should be communicated and the stepped psychological and pharmacological management that follows.
- Recall the prognosis of a functional seizure disorder and what changes when it coexists with genuine epilepsy.

The material this chapter leans on but does not own is worth naming so that revision goes to the right place. The general construct of a functional neurological disorder, within which the functional seizure sits, is developed in the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes. The mood-stabiliser lens on valproate, carbamazepine and lamotrigine, their bipolar indications and monitoring, belongs to the Mood disorders: pharmacotherapy notes, while this chapter keeps only the antiepileptic lens. The full adverse-effect profile of clozapine, beyond the seizure-threshold angle, is the Schizophrenia: management, antipsychotics and adverse effects notes. Delirium in its own right, with its formal assessment and prevention, is the Stroke, TBI, delirium and focal brain syndromes notes, and the full anatomy of the limbic system is the Functional Neuroanatomy and Neurophysiology notes. Cross-refer to these by name; do not try to carry their numbers in this chapter's memory.

IN INDIA

Several of these topics carry a distinctly Indian answer. Valproate in a woman of childbearing potential is a harder decision where reliable contraception and pre-conception planning are less assured, and the drug remains widely prescribed on grounds of cost and availability. Phenobarbital continues in first-line use in many resource-limited settings for the same reasons, which shapes both the seizure control the psychiatrist inherits and the behavioural and cognitive burden that comes with it. A functional seizure disorder often presents late and heavily stigmatised, sometimes framed by the family in spiritual rather than medical terms, so the way the diagnosis is communicated matters even more than in the settings where the evidence base was written. Video-electroencephalography sits in tertiary centres alone, so much of the peri-ictal placement and the exclusion of a functional event is a clinical judgement made from a careful history rather than from a recording.

45.8 The master map, and walking into the examination

One table to sweep the whole field before you turn the paper. The map below compresses the chapter to a single view: each domain, the discrimination the examiner turns on, and the withheld targets you should be able to produce. Read it as a self-test: if any recall target is blank, the owning topic is where to go, not this cell.

DOMAIN	THE DISCRIMINATION THE EXAMINER TURNS ON	RECALL TARGETS (SUPPLY THE ANSWERS)
Classification and localisation	focal versus generalised onset, and which localisation is psychiatrically loud	the temporal-lobe and frontal-lobe behavioural signatures and their non-epileptic mimics
The peri-ictal axis	place any symptom by its timing to the seizure	the temporal windows and which syndrome occupies each
The psychoses of epilepsy	ictal versus postictal versus chronic interictal, plus forced normalisation	the lucid interval, the typical durations, and the interictal record state for each type
Epileptic versus primary psychosis	what an epileptic psychosis preserves that a primary disorder loses	the three features that mark it as epileptic rather than primary
Mood, anxiety and suicide	the epilepsy-intrinsic story versus the drug-attributed signal	the epilepsy-specific screen, the suicide-risk magnitude, and the regulator warning
Personality, behaviour and cognition	disease versus drug attribution	the interictal change, the eponymous cluster, and the drug cognitive profile
Functional events and the drug interface	functional event versus true seizure, and mood-stabiliser versus antiepileptic lens	the differentiating semiology, the confirmatory investigations, and where the mood-stabiliser evidence lives

A revision map of the chapter: the withheld recall targets are owned and developed by the individual topics and are named here only to be tested, never re-taught.

When the question actually arrives, run the three questions of the master frame in order and let them route the answer: situate the disturbance in time, decide whether the disease, the drug, or an independent disorder is responsible, and name the imitator you have excluded to reach the label. A candidate who reasons this way in the viva, rather than reciting a syndrome from the middle, shows exactly the discrimination the examination is built to reward, and the individual facts fall into place around it because they were learned as answers rather than as a list. The chapter is long, but the skill it teaches is short: timestamp, attribute, exclude.

When	Whose	What else	Then treat
PLACE IT ON THE PERI-ICTAL AXIS	THE DISEASE, THE DRUG, OR INDEPENDENT	THE IMITATOR TO EXCLUDE	SAFETY, AND THE SEIZURE THRESHOLD

References

The canonical sources this chapter is grounded in. Every fact-bearing specific in the notes traces to one of these; where a primary paper is named in the text, it is cited inline.

1. David AS, Fleminger S, Kopelman MD, Lovestone S, Mellers JDC. Lishman s Organic Psychiatry: A Textbook of Neuropsychiatry. 4th ed. Wiley-Blackwell.
2. Boland R, Verduin ML, eds. Kaplan and Sadock s Comprehensive Textbook of Psychiatry. 11th ed. Wolters Kluwer; 2024.
3. Kaufman DM, Geyer HL, Milstein MJ. Kaufman s Clinical Neurology for Psychiatrists. 9th ed. Elsevier.
4. Taylor DM, Barnes TRE, Young AH. The Maudsley Prescribing Guidelines in Psychiatry. 15th ed. Wiley-Blackwell.
5. Standard EEG references (Niedermeyer s Electroencephalography; clinical EEG atlases).
6. Martin A, Bloch MH, Volkmar FR, eds. Lewis s Child and Adolescent Psychiatry. 5th ed. Wolters Kluwer.
7. Brunton LL, et al, eds. Goodman and Gilman s The Pharmacological Basis of Therapeutics. McGraw-Hill.
8. American Psychiatric Association. Diagnostic and Statistical Manual of Mental Disorders, 5th ed, Text Revision (DSM-5-TR). APA; 2022.
9. World Health Organization. ICD-11 Clinical Descriptions and Diagnostic Requirements (CDDR). WHO.
10. Johnstone EC, et al, eds. Companion to Psychiatric Studies. Churchill Livingstone.

Additional primary sources cited inline: Dulcan's Textbook of Child and Adolescent Psychiatry, Panksepp Affective Neuroscience, RC Jiloha Forensic Psychiatry.

PART II

Neurology Foundations for Psychiatrists

The neurology a psychiatrist actually needs: where a lesion puts itself, the headache that is not migraine, and the sleep disorder hiding behind the mood complaint.

01 The bedside examination, what it localises, and the investigations it earns

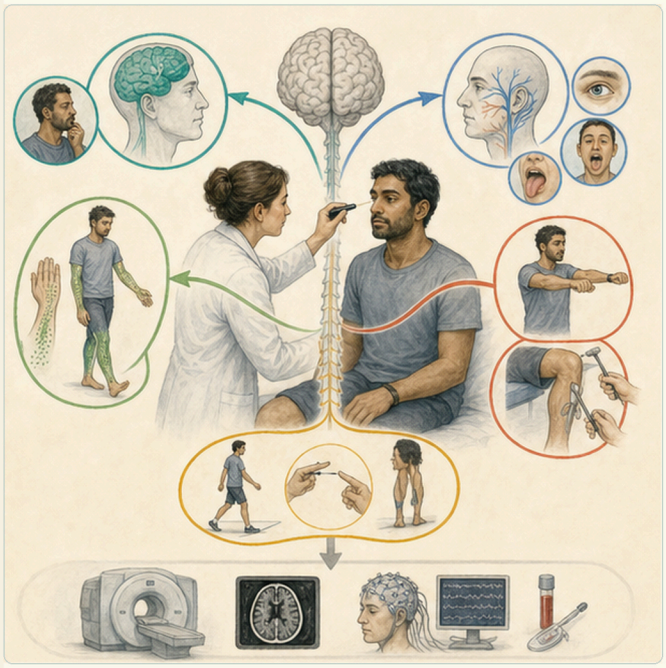
02 Headache and the psychiatric interface

03 Sleep: the architecture and its disorders

04 Consolidate and apply

05 References

VISUAL 2.1



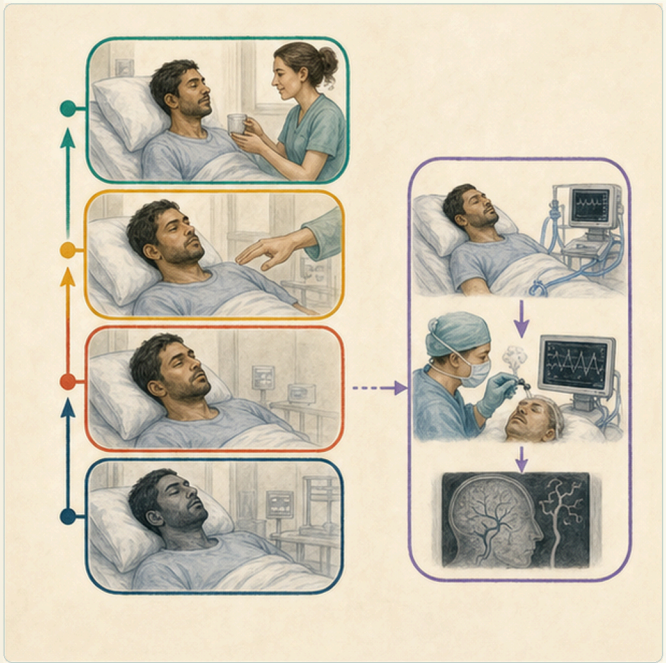
The neurological examination earns localisation

- Higher functions
- Cranial nerves
- Motor and reflexes
- Sensory
- Cerebellar and gait
- Targeted investigations

The examination is organised to localise the lesion first and to justify targeted investigation second.

Source lock: day-31.final.html#sec-t02

VISUAL 2.2



Disorders of consciousness as observable states

- Coma
- Unresponsive wakefulness
- Minimally conscious state
- Emergence
- Brain-death determination

Consciousness states are distinguished by arousal and reproducible purposeful behaviour; brain-death determination is a separate clinical process.

Source lock: day-31.final.html#sec-t07

VISUAL 2.3



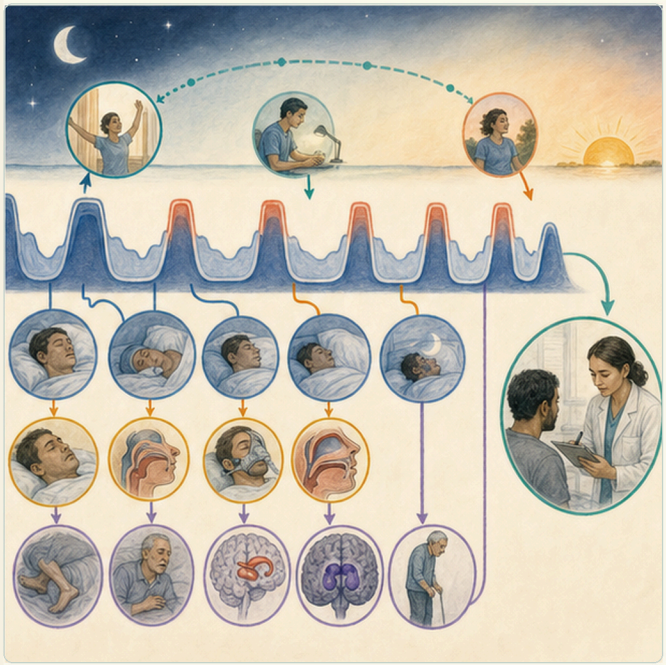
Headache at the psychiatric interface

- Primary headache
- Red flags
- Psychiatric comorbidity
- Medication overuse
- Prevention

Headache assessment holds primary headache, secondary red flags, psychiatric comorbidity, medication overuse and prevention in one formulation.

Source lock: day-31.final.html#sec-t08

VISUAL 2.4



Sleep architecture connects night to diagnosis

- NREM
- REM
- Circadian rhythm
- Breathing
- Movement and neurodegeneration

Sleep complaints become clinically legible when architecture, circadian timing, breathing and movement phenomena are examined together.

Source lock: day-31.final.html#sec-t12

PAPER IV

Neurology foundations for psychiatrists

Three bands . one complete notes document

The bedside examination, what it localises, and the investigations it earns

1 · Neurology for the psychiatrist: what this chapter owns, and what it points to

Orientation is the gatekeeper of the neurological assessment in psychiatry. It establishes whether the patient can take part in a meaningful interview, follow the examination, retain the immediate purpose of the encounter and give information that can safely be interpreted. This is not a decorative opening question. It is the first safeguard against mistaking a disorder of attention, arousal, memory, language, perception or consciousness for a purely psychiatric account of behaviour.

For the psychiatrist, orientation is best understood as a bedside judgement about the patient's current relationship to self, place, time and situation. The answer matters less as a recital of isolated facts than as evidence about how the patient is organising present experience. A person may state a fact correctly yet have reached it by guesswork, rehearsal, cueing from relatives or a fragmentary understanding of the encounter. Conversely, an anxious, frightened, poorly educated or sensory-impaired patient may hesitate without being globally disoriented. The examiner therefore asks, listens, observes and then decides how much weight the remainder of the interview can bear.

1.1 What orientation tests at the bedside

Orientation is the capacity to locate oneself coherently in the immediate world. In clinical work it is not a single faculty. It is the visible product of sufficient wakefulness, sustained attention, registration of current information, access to autobiographical knowledge, language comprehension and the ability to communicate an answer. The patient must also understand the social meaning of the question. A reply to a question about place is only useful if the patient knows that the question concerns present location rather than domicile, birthplace, workplace or an imagined setting.

The practical value of orientation lies in its economy. It gives the examiner an early sample of contact, comprehension and reliability before the history becomes detailed. It also changes the stance of the interview. When orientation is sound, a psychiatrist can generally proceed to the higher-function, affective, psychotic and risk assessment while remaining alert to focal clues. When orientation is doubtful, the task becomes more basic: establish the level and fluctuation of awareness, obtain collateral information, examine for acute medical or neurological contributors and avoid treating every unusual statement as a belief held with normal alertness and cognitive access.

Orientation is related to, but is not synonymous with, **attention**. A person can give a practised answer to routine orientation questions while being unable to hold the thread of conversation. Equally, an attentive person may be uncertain about an unfamiliar environment without having a global disturbance. The difference is clinically important because attention tells the examiner whether information is being taken in now, while orientation tells the examiner whether the patient can place current experience into a usable frame. A brief answer never replaces observation across the interview.

Nor is orientation simply a test of factual knowledge. A patient may know general public facts yet be unable to identify the current setting, or may fail a conventional question because the setting has been changed, renamed or poorly explained. The clinician should seek the level of practical orientation needed for the immediate decision: Can the patient understand who is speaking, where care is occurring, what has happened and what is being proposed? This makes orientation clinically useful in unfamiliar wards, emergency settings, homes, clinics and remote encounters alike. It also prevents needless confrontation over a detail that has little bearing on the patient's present capacity to engage.

RULE **Treat orientation as evidence about the reliability of the encounter, not as a pass-or-fail recital.** Interpret each answer with its route to the answer, the patient's alertness, language, sensory access and the consistency of the surrounding interview.

1.2 The domains and their clinical meaning

Inquiry commonly moves through orientation to person, place, time and situation. These domains should be explored conversationally, using questions suited to the patient's language, culture and immediate context. The aim is not to catch the patient out. It is to determine whether the internal account is coherent, whether uncertainty is acknowledged appropriately and whether an answer remains stable when the question is approached from another angle.

- **Orientation to person** concerns the patient's sense of personal identity and recognition of relevant people. A difficulty here can reflect a disturbance of self-knowledge, memory, recognition, language, psychosis or the meaning of the social interaction. It should prompt careful exploration rather than an immediate assumption about severity.
- **Orientation to place** asks whether the patient can identify the present setting in a way that fits the encounter. Confusion may be revealed not only by a wrong location but by a plausible location that does not explain why the patient is there.
- **Orientation to time** samples the patient's grasp of temporal context. The useful clinical question is whether the patient can situate the present moment sufficiently to make sense of events, medication, admission, sleep and symptom onset, not whether every calendar detail can be recited.
- **Orientation to situation** asks whether the patient understands what is happening and why professional contact is occurring. It connects cognitive orientation to judgement, insight and the immediate practical demands of care.

These domains should not be forced into a false hierarchy. Their pattern may help the clinician describe the problem, but no domain by itself locates a lesion or establishes a diagnosis. The anatomical framework belongs in the Functional Neuroanatomy and Neurophysiology notes, and the detailed neurological examination belongs in the companion examination section of this chapter. At this bedside stage, the psychiatrist's duty is to recognise when the pattern is discordant with the proposed psychiatric formulation and to make that discordance visible in the record.

MNEMONIC

PERSON - PLACE - TIME - SITUATION

Use this as a quiet mental sweep, not a script. After each domain, ask whether the answer is coherent, independently generated and consistent with the patient's alertness and the practical purpose of the assessment.

1.3 How to ask without creating an artefact

Good orientation testing begins before the formal question. Notice whether the patient turns toward the examiner, tracks the topic, responds after a reasonable pause, asks for clarification when necessary and recognises why questions are being asked. Observe whether the patient is hearing adequately, can see environmental cues, has access to a familiar language and is being overwhelmed by pain, breathlessness, fear, intoxication, fatigue or a crowded clinical setting. These conditions can alter performance without proving a disorder of orientation.

Questions should be simple, direct and open enough to show the patient's own frame of reference. Ask for clarification rather than supplying the answer. If an answer is uncertain, revisit it later in a naturally different form. This distinguishes a transient lapse, a misunderstood question and a stable inconsistency. It also keeps the examination humane. A patient who feels tested may become guarded, oppositional or ashamed, especially when already frightened by hospitalisation. Such reactions are data about the interaction, not proof that cognition is intact or impaired.

Collateral information is part of this task, not an afterthought. Relatives, nursing staff and records can clarify the patient's usual baseline, recent change, fluctuations, exposure to substances or medicines, and the circumstances of presentation. Their account should support rather than replace direct examination. The psychiatrist should record the source and distinguish observed behaviour from reported information. When the patient cannot participate reliably, the history becomes a reconstruction, and the limits of that reconstruction must be stated plainly.

TRAP **A correct answer is not automatically an intact function.** Rehearsed responses, environmental prompting and preserved fragments of knowledge can conceal impaired attention or a fluctuating state; judge the answer in the flow of the interview.

The examiner must reduce avoidable disadvantage before interpreting failure. Use the patient's preferred language where feasible, arrange communication support where needed and

make sensory aids available. Consider whether unfamiliarity with formal institutions, literacy demands or cultural assumptions have changed the question's meaning. This is not a lowering of clinical standards. It is how the examination becomes specific to the function it claims to assess rather than to the patient's social comfort with the examiner.

1.4 Separating orientation from the psychiatric mental state

Disorientation can coexist with depression, mania, anxiety, psychosis, dissociation, substance-related states and behavioural disturbance, but it must not be absorbed into any of them without thought. The content of a statement and the conditions under which it is produced are separate clinical questions. A patient may voice an implausible explanation because they hold a delusional belief, because they are confused about their surroundings, because language has failed, because they are attempting to fill a gap in memory, or because they do not understand the question. These possibilities can overlap. Their distinction determines whether the interview is primarily a phenomenological assessment, an urgent medical assessment or both.

Insight is also not interchangeable with orientation. A person may accurately identify place and purpose while rejecting illness, declining treatment or interpreting events through a fixed belief system. Conversely, a person may express apparent acceptance of illness while being unable to describe why they are in the setting or what has just occurred. Psychiatry often rewards a richly elaborated account, but elaboration should not be mistaken for reliability. The more fluent the narrative, the more important it is to check whether it remains anchored to the immediate context.

Dissociation deserves the same discipline. Experiences of unreality, detachment, altered identity or gaps in recall can be clinically compelling, yet they do not remove the need to examine alertness, attention and orientation. The psychiatrist should neither dismiss a dissociative account as feigned confusion nor label a confused patient as dissociative because the history sounds psychologically meaningful. The relevant phenomenology is covered in the Anxiety, OCD and trauma-related disorders notes and the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes; at the neurological bedside, the immediate job is to establish what the patient can currently know, hold and use.

Formal cognitive screens and more extensive cognitive assessment are not substitutes for this bedside reasoning. Their selection, scoring and interpretation belong in the Psychometry, Assessment and Descriptive Psychopathology notes. Here, an abnormal orientation response should lead to a richer account of the patient's functioning, not to a score-driven conclusion divorced from arousal, communication and collateral evidence.

> **PITFALL**
>
> Do not call an implausible answer "psychotic" before establishing that the patient understood the question, could attend to it and can place the encounter in present reality. Content without cognitive context is an unsafe shortcut.

1.5 Recognising when orientation changes the clinical priority

Fluctuation is often more informative than a single failed item. A patient who is coherent in one moment and unable to sustain the same frame of reference later requires repeated observation and a timeline from those who have been present. Variation in alertness, attention, speech, behaviour or orientation should move the clinician away from a one-off mental-state label and toward an account of changing brain function. This does not itself establish cause, but it changes urgency and broadens the differential.

Do not confuse apparent improvement with reliable recovery. A patient may seem clearer after rest, reassurance or environmental simplification, yet still be unable to maintain the new level of engagement when demands increase. In the same way, an apparently impaired response may improve once pain is relieved, a familiar person is present or communication is made accessible. Serial assessment therefore asks a simple but powerful question: is the patient's orientation stable enough for the decision now required? The answer can be qualified. It need not become a claim of normality before the psychiatrist can communicate safely with the wider team.

Where orientation is impaired or cannot be confidently assessed, the psychiatrist should first make the encounter safe and information-rich. Confirm immediate physiological and behavioural safety through the appropriate clinical team; seek collateral; review the temporal relation to illness, injury, substances and treatment; and repeat the examination as the state permits. The detail of the physical neurological examination is taught elsewhere in this chapter. Investigation choices and exam-to-imaging reasoning are likewise separate from the imaging modalities described in the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes. The important orientation lesson is that a psychiatric formulation is provisional when its cognitive platform is unstable.

The same caution applies to severe agitation or withdrawal. The observable syndrome may make a conventional interview impossible, but that is a reason to adjust method, not to abandon neurological thinking. Describe whether the patient can be engaged, whether responses are purposeful, whether orientation can be sampled at all and what prevents a reliable assessment. This preserves clinical uncertainty honestly and helps the next clinician detect change rather than inherit a misleading categorical label.

> **GOOD TO REMEMBER**
>
> When orientation is doubtful, document what the patient did, not only what they failed. The route of response, consistency, environmental cues, barriers to testing and collateral baseline often guide the next assessment more than a bare statement that the patient was "disoriented".

1.6 Recording an orientation assessment that another clinician can use

A useful note identifies the domains explored, the patient's responses in plain language, the degree of confidence in those responses and the conditions of testing. It should show whether information was volunteered, prompted or inferred from behaviour. It should also name important limitations: reduced arousal, inattention, language mismatch, hearing or visual

difficulty, distress, inability to cooperate, sedating treatment, or missing collateral. Avoid replacing these observations with a global adjective that conceals uncertainty.

Record the temporal story. Is the presentation described as new, longstanding, episodic, progressive, sudden or variable? Was the patient reportedly functioning differently earlier? Does the patient's account agree with observed behaviour and collateral reports? Such language does not make a diagnosis, but it creates a defensible bridge from the psychiatric interview to medical review. It also supports serial assessment. A later clinician should be able to tell whether the person's ability to orient has changed, or whether the apparent change merely reflects a different question, a different setting or a different level of support.

In handover, communicate the practical implication. State whether the patient can give a reliable account, participate in consent discussions, report symptoms accurately and be left to self-report risk. If uncertainty remains, say what needs reassessment and under what conditions. This is particularly important when deciding how much meaning to assign to mood, psychotic content, refusal, apparent non-cooperation or a request to leave. Orientation is not a bureaucratic field in the mental state examination. It is a declaration of the evidential quality of the rest of it.

BEDSIDE OBSERVATION	WHAT TO CLARIFY	HOW TO RECORD THE IMPLICATION
Confident response	Whether it is independently generated and fits the conversation	State the context supporting reliability
Uncertain response	Question comprehension, sensory access, anxiety and familiarity with the setting	Describe the barrier before inferring impairment
Inconsistent response	Attention, fluctuation, cueing and collateral baseline	Record variation and need for serial review
Unable to engage	Arousal, behaviour, communication and immediate safety	State that orientation cannot be reliably assessed and why

1.7 The neurological lens for psychiatry

Orientation teaches a habit that extends through the whole psychiatric examination: first establish the conditions under which a symptom can be known. The question is not merely "What does the patient say?" It is "What current brain state, communicative capacity and environmental context made that statement possible?" This protects against premature closure in both directions. It prevents the clinician from medicalising every unusual experience, but it also prevents a psychiatric diagnosis from becoming a reason not to notice altered cognition.

That habit should remain proportionate. A patient with a coherent interview, stable orientation and no discordant neurological clues does not need a theatrical neurological workup merely because they have a psychiatric syndrome. Conversely, a polished psychiatric history should not reassure the clinician when it is produced on a platform of impaired attention, changing awareness or unreliable orientation. The subsequent examination, localisation reasoning and imaging decision are covered in their designated sections. Sleep phenomena require the parallel

neurological and psychiatric framing in the Eating and sleep-wake disorders notes. Electroencephalographic questions belong with the Epilepsy and psychiatry notes.

Self	Place	Time	Situation
IDENTITY AND RELEVANT PEOPLE	CURRENT SETTING	PRESENT TEMPORAL CONTEXT	PURPOSE OF THE ENCOUNTER

Orientation is not a memory quiz: it establishes whether the psychiatric interview rests on a reliable account of present reality.

2 · The neurological examination relevant to psychiatry (higher functions, cranial nerves, motor, sensory, reflexes, cerebellar, gait)

The neurological examination is not an ornamental appendix to the psychiatric assessment. It is the bedside method that separates impaired cognition from impaired access to cognition, medication effects from primary motor disease, affective expression from corticobulbar dysfunction, and functional symptoms from deficits that obey neuroanatomy. Marks are earned by showing an orderly examination, describing the sign precisely, and stating what the pattern localises.

The safest routine is a **fixed sequence**: mental status, cranial nerves, motor system, reflexes, sensation, cerebellar system, and gait. Mental status comes first because impaired attention, comprehension, or language can make every later response unreliable. The accompanying figure should be read as a working bedside route, not as a checklist to recite without examining the patient.

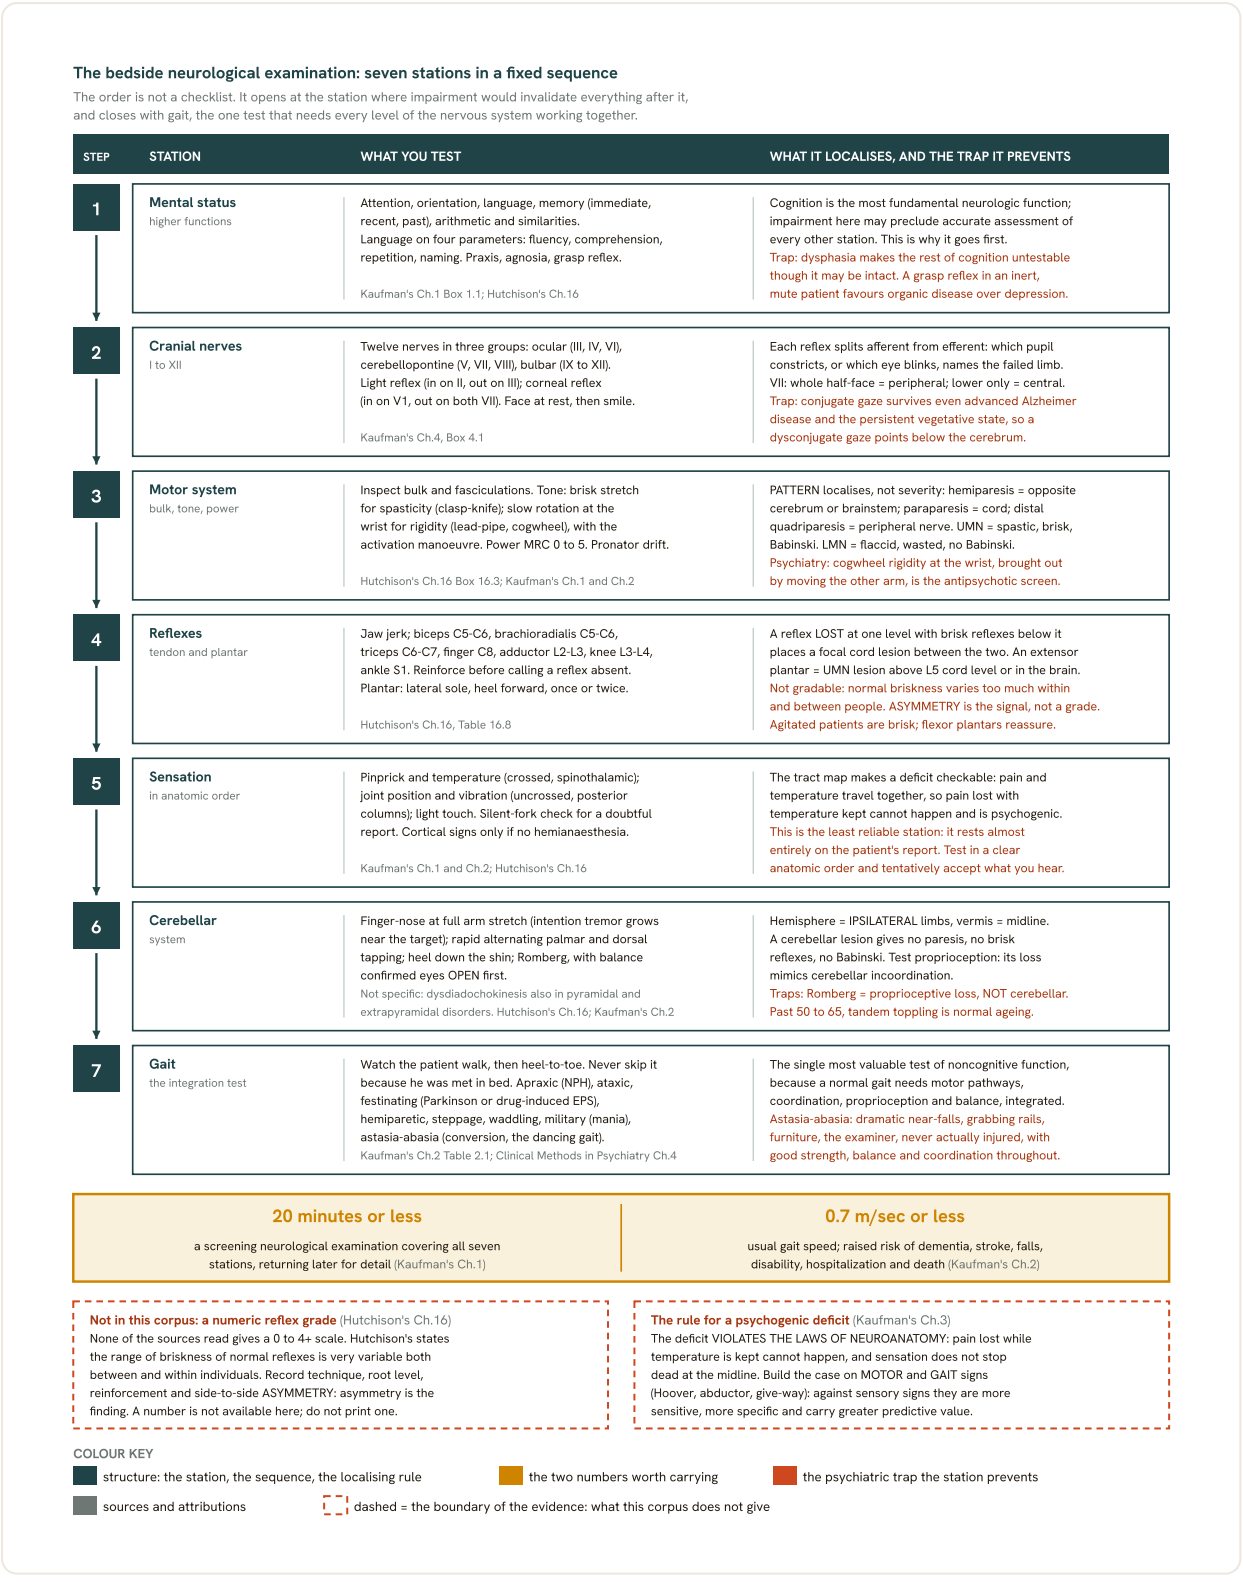


Fig 31.1 **The bedside neurological examination as a structured sequence.** The seven stations run in the order given by Kaufman's Clinical Neurology for Psychiatrists (Box 1.1): mental status, cranial nerves, motor system, reflexes, sensation, cerebellar system and gait. The order is itself the teaching: mental status leads because cognitive impairment may preclude accurate assessment of every other station, and gait closes because it needs every level intact and integrated, which is what makes it the single most valuable assessment of noncognitive nervous system function. Each station earns its place in a psychiatric assessment by the misreading it prevents. Attributions per box (Kaufman's Clinical Neurology for Psychiatrists, Hutchison's Clinical Methods, Clinical Methods in Psychiatry (AIIMS)).

MIND	PATTERN	ASYMMETRY	GAIT
VALIDATES	BEATS SEVERITY	LOCALISES	INTEGRATES

2.1 Conduct, sequence, and formulation

A screening examination can usually be completed in **twenty minutes or less**, provided the examiner uses the same route every time and returns for detail where an abnormality appears. Do not abandon the sequence because an obvious deficit is found. A florid tremor does not exclude pyramidal weakness, and a psychiatric history does not make an objective cranial nerve sign less important.

MNEMONIC

Mind, Cranium, Movement, Reflex, Sensation, Coordination, Walking

Use this as a prompt for the sequence: **mind** before instructions, **cranium** for cranial nerves, **movement** for bulk, tone and power, then **reflex**, **sensation**, **coordination**, and **walking**.

The examination should answer what the patient experiences, what the examiner can demonstrate, where the dysfunction lies, and what processes could produce it. “The leg feels dead” is a symptom; a reproducible pyramidal pattern, reflex asymmetry, or inconsistency during synergistic movement is a sign. Move from description to localisation before diagnosis. Imaging follows that reasoning; full pathway anatomy belongs to the Functional Neuroanatomy and Neurophysiology notes.

RULE **Examine in sequence, interpret in patterns, and localise before naming a disease.** A single brisk reflex or uncertain sensory boundary is weak evidence; convergent findings across independent parts of the examination are much stronger.

Before testing, observe the patient entering, sitting, speaking, and reaching. Explain each manoeuvre, compare sides, and retest unexpected findings differently. If arousal is inadequate, higher-function testing cannot be interpreted. If language or hearing is impaired, alter the method rather than labelling the patient uncooperative.

2.2 Higher functions at the bedside

The bedside mental-status component samples attention and cooperation, orientation, language, memory for immediate, recent and remote events, and higher intellectual functions such as calculation and abstraction. It begins during spontaneous conversation. Note whether the patient sustains the task, understands its purpose, initiates a response, monitors errors, and persists. These observations often explain a poor test performance more accurately than the final answer alone.

Bedside language examination samples fluency, comprehension, repetition, and naming. Listen for effort, phrase length, grammatical structure, paraphasic errors, empty fluent output, and awareness of mistakes. Test comprehension with progressively less contextual commands, repetition with material that cannot be guessed from context, and naming with familiar objects and their parts. Naming is vulnerable across dysphasic syndromes. Severe comprehension impairment can invalidate apparent memory, calculation, praxis, and executive testing even when those capacities remain intact. In psychiatric practice this prevents fluent but meaningless speech

being mistaken for formal thought disorder, and non-fluent output being mistaken for mutism, depression, or deliberate refusal.

Apraxia is inability to perform a learned act despite sufficient motor and sensory capacity. Ask the patient to mime tool use, wave, salute, or copy demonstrated hand postures. Confirm that weakness, sensory loss, poor comprehension, and inattention do not explain failure. A multistep command is broader than a pure praxis test because it also recruits language, memory, sequencing, and executive control.

Agnosia is modality-specific failure of recognition despite preserved basic sensation and relevant semantic knowledge. A patient who cannot identify an object by sight may identify it by touch and may still describe its purpose. The comparison between modalities is the point. Similarly, cortical sensory testing asks whether a normally felt stimulus can be integrated and recognised, rather than whether it is detected at all.

Formal screens and neuropsychological batteries may quantify impairment, but administration and scoring belong to the Psychometry, Assessment and Descriptive Psychopathology notes. At this bedside stage, the psychiatrist should state the function tested, control for confounds, describe the error, and say whether it suggests a focal language or praxis disorder, a diffuse cognitive problem, or unreliable access to otherwise preserved function.

2.3 Cranial nerves: a psychiatry-focused sweep

Test smell one nostril at a time with a familiar, innocuous aromatic substance. Irritative substances are unsuitable because nasal trigeminal endings can detect them despite olfactory loss. Smell complaints matter because self-report is unreliable and anatomically impossible performance may reveal invalid responding. On a four-alternative forced-choice task, performance **below twenty-five per cent** is below chance and cannot be explained by genuine anosmia alone.

For the optic nerve, assess acuity, fields by confrontation, and the fundi. Examine each eye independently before comparing them; binocular field testing can miss a chiasmal pattern. In the **pupillary light reflex**, the optic nerve carries the incoming signal and the oculomotor nerve carries the constrictor response. Light in either eye should constrict both pupils. Failure after stimulation from a particular eye, with preserved bilateral responses from the other, indicates an afferent defect. Preserved consensual constriction with failure of a particular pupil to constrict indicates an efferent defect on that side.

Inspect lids, pupils, resting eye position, and conjugacy, then test pursuit while asking about diplopia. Oculomotor impairment combines ptosis, a dilated pupil, and outward deviation. Abducens impairment produces inward deviation without lid or pupil findings. Trochlear injury may produce compensatory head tilt away from the affected side. Dysconjugacy should not be attributed to cortical dementia: preserved conjugacy despite advanced cerebral disease directs abnormal findings toward the brainstem, peripheral nerve course, or neuromuscular transmission.

Test facial sensation in the forehead, malar region, and jaw distribution, preserving the anatomical boundaries. The angle of the jaw is not supplied by the trigeminal territory, and complete facial loss that stops sharply at the hairline is anatomically suspect. In the **corneal reflex**, the ophthalmic trigeminal division is afferent and facial nerves provide the bilateral blink. Absent blinking after stimulation from a particular cornea, with bilateral blinking from the other, points to an afferent lesion; failure of a particular eyelid to close whichever cornea is stimulated points to an ipsilateral facial efferent lesion.

Observe the face at rest and during forehead wrinkling, eye closure, and smiling. A peripheral lesion weakens the ipsilateral upper and lower face; a central corticobulbar lesion predominantly weakens the contralateral lower face because upper facial movement has bilateral cortical input. Hearing can be screened by whispered material presented to each ear separately. Unilateral central lesions do not produce deafness because auditory representation is bilateral; apparent unilateral deafness beside hemiparesis deserves anatomical review.

For bulbar function, listen to spontaneous speech and ask for repeated lingual, labial, and pharyngeal syllables, then inspect palatal elevation during phonation. Lower motor bulbar and corticobulbar syndromes share dysarthria and dysphagia, so these do not make the distinction. Hypoactive gag and jaw responses, respiratory involvement, and absent emotional lability favour bulbar dysfunction; brisk gag and jaw responses with emotional lability and cognitive impairment favour pseudobulbar dysfunction. In **pseudobulbar affect**, crying or laughter is involuntary and dissociated from the underlying mood; true depression may nevertheless coexist. Finally inspect the tongue at rest and on protrusion: a unilateral hypoglossal lesion produces wasting and deviation toward the weak side.

2.4 Motor system: bulk, tone, power, and pattern

Inspect before touching. Compare muscle bulk, posture, spontaneous movement, tremor, and fasciculation. Marked wasting suggests denervation and a lower motor pathway problem; upper motor lesions produce much less wasting, and later. Apparent hollowing in a thin or older person is not enough if strength is preserved. Fasciculations are visible twitches of individual motor units in resting muscle. They gain meaning from distribution, weakness, and wasting, not from their mere presence.

Assess tone with the patient relaxed and by passive movement at different speeds. Spasticity is velocity-dependent: brisk stretch produces a catch and may then yield. Extrapyramidal rigidity resists movement through the range and affects opposing muscle groups; it may feel uniform or intermittently ratcheted by tremor. Test the wrist with passive flexion, extension, and rotation, and use contralateral voluntary movement to reveal subtle rigidity. This is particularly useful when screening a patient receiving dopamine-blocking medication. Active resistance in every direction may instead be paratonia, voluntary guarding, pain, or catatonic negativism, and the quality and context must be described.

Power is tested isometrically by opposing the intended action. Compare sides and proximal with distal groups, stabilise the joint, and prevent substitution. Most groups require separate testing in

each limb. Mild plantar-flexion weakness may be clearer during toe walking and dorsiflexion weakness during heel walking. The **Medical Research Council scale** has **six grades from zero to five**:

- Grade 0 is no contraction and grade 1 is a trace.
- Grade 2 is movement with gravity eliminated and grade 3 is movement against gravity.
- Grade 4 is movement against resistance and grade 5 is normal power.

Record the movement tested, not only a global grade. The scale compresses much clinically important weakness into the resistance grade and works poorly for small distal movements where gravity contributes little. Functional capacity therefore remains essential.

FEATURE	UPPER MOTOR PATHWAY PATTERN	LOWER MOTOR PATHWAY PATTERN
Weakness	Pyramidal distribution	Myotomal, nerve, or distal pattern
Bulk	Relatively preserved early	Marked wasting with denervation
Tone	Spastic	Flaccid
Tendon reflexes	Pathologically brisk	Reduced or absent
Plantar response	Extensor response may be elicited	No extensor response from the lower motor lesion itself

Pattern localises better than absolute severity. Hemiparesis involving lower face, arm, and leg suggests the contralateral cerebral hemisphere or brainstem; paraparesis directs attention to the spinal cord; distal weakness affecting all limbs favours the peripheral nervous system. For subtle pyramidal weakness, ask the patient to extend and supinate the arms with eyes closed. Downward drift with pronation is abnormal. This manoeuvre exposes a coherent motor pattern, whereas an isolated complaint of heaviness does not.

2.5 Reflexes: objective signs and level logic

Tendon reflexes are monosynaptic stretch responses. They test an intact sensory afferent, spinal connection, motor efferent, muscle, and descending modulation. A missing response is therefore not automatically a motor lesion; sensory neuropathy or radicular disease may interrupt the afferent limb. Position the limb comfortably, ensure relaxation, strike the tendon briskly, and watch as well as feel the contraction. Compare sides at matched levels.

REFLEX	PRINCIPAL ROOTS	PRACTICAL TECHNIQUE
Biceps	C5-C6	Place a finger on the tendon and strike the finger
Brachioradialis	C5-C6	Place a finger over the tendon region and strike
Triceps	C6-C7	Support the relaxed arm and strike the tendon
Finger	C8	Elicit with the hand relaxed
Adductor	L2-L3	Observe adductor contraction
Knee	L3-L4	Support the flexed leg and strike the patellar tendon
Ankle	S1	Hold the foot and strike the calcaneal tendon

Normal briskness varies widely, especially with anxiety and incomplete relaxation. Reinforce an apparently absent lower-limb reflex by asking the patient to interlock and pull the fingers at the moment of the strike; reinforce an upper-limb response with contralateral gripping. Do not diagnose neuropathy until reinforcement has been attempted. Conversely, do not diagnose an upper motor lesion from briskness alone. Asymmetry, spread, clonus, tone, weakness pattern, and plantar response supply the context.

The jaw jerk is useful when dysarthria, dysphagia, emotional lability, or widespread spasticity raises a corticobulbar question. Place a finger across the relaxed open chin and tap it lightly. A soft rebound is expected; a depressed response supports lower motor bulbar dysfunction, while a forceful response supports corticobulbar dysfunction. Segmental localisation uses a lost reflex at a given level with pathologically brisk reflexes below it to identify a focal cord level.

To elicit the plantar response, warn the patient, support the foot, and draw a blunt disposable stimulus firmly along the lateral sole from the heel toward the base of the little toe. The minimum extensor response is upward movement of the great toe; fanning may accompany it but is not required. The **Babinski sign** is described as present or absent, not “positive” or “negative.” In an agitated patient with increased tone and brisk reflexes, flexor plantar responses are reassuring. When definite spasticity, pyramidal weakness, pathologically brisk reflexes, and clonus already converge, the plantar response adds less; when the other signs are equivocal, it may decide whether a central lesion is present.

2.6 Sensory examination: tract logic with disciplined scepticism

The sensory examination is the least objective part of the neurological examination because it depends almost entirely on the patient's report. The answer is not to dismiss it. Define the question, demonstrate the stimulus, keep timing and force unpredictable, compare homologous sites, proceed in an anatomical order, and tentatively accept the report while looking for reproducibility and functional consequences.

Test light touch, pain, temperature when indicated, vibration, and joint position. Pain and temperature travel together in the lateral spinothalamic system after crossing near entry to the cord. Position, vibration, and discriminative cortical functions depend on posterior-column input ascending ipsilaterally before crossing in the medulla. Light touch also has anterior spinothalamic representation. This arrangement makes selective preservation of temperature with claimed complete loss of pain anatomically incoherent, while dissociation between pain and position can be anatomically meaningful.

For position sense, close the patient's eyes, hold the sides rather than the top and bottom of a distal phalanx, and move it slightly upward or downward in random order. Side-grasping removes pressure clues. If distal perception is impaired, move proximally and use larger excursions. For vibration, apply a noiselessly activated low-frequency fork to distal bony points and move proximally if sensation is reduced. If reports appear unreliable, randomly apply the fork vibrating or silent while the patient's eyes remain closed. The concealed comparison tests validity without confrontation.

Cortical sensory testing is meaningful only when primary sensation is adequate. For extinction, touch either side alone and then both simultaneously; detection of each alone but omission of a side during bilateral stimulation indicates sensory inattention. For tactile recognition, ask the patient with eyes closed to identify a familiar object that can be named on sight. For graphesthesia, compare recognition of writing on the palm with eyes open and closed. These tests require attention, language, and cooperation.

PITFALL

Do not diagnose a functional sensory deficit merely because the map looks odd. Repeat the examination, check comprehension and attention, test tract-linked modalities, and look for objective motor, reflex, cerebellar, and gait evidence. Sensory signs carry less predictive weight than positive motor and gait signs.

2.7 Cerebellar examination, stance, and gait

Cerebellar hemispheres coordinate ipsilateral limbs, while the vermis coordinates axial structures. Cerebellar lesions cause incoordination rather than primary weakness, pathologically brisk reflexes, or an extensor plantar response. Examine speech, eyes, upper and lower limbs, stance, and gait, while testing proprioception because sensory ataxia can imitate cerebellar incoordination.

For the finger-nose test, place your finger far enough away to require a full reach and ask the patient to alternate between it and the nose. Cerebellar dysmetria produces an irregular trajectory and an action tremor whose amplitude increases near the target. Rapid alternating movement is tested by repeated alternating palm and dorsum contact on the same spot. Uneven force, irregular rhythm, and loss of alternation support ipsilateral incoordination, but this finding is not specific. For the heel-shin test, the supine patient places the lifted heel below the opposite knee and runs it down the shin. Tremulous placement and a wavering descent indicate incoordination.

Romberg's test begins with the feet together and eyes open. Only after confirming that the patient can stand in that condition should the eyes be closed, with the examiner ready to catch the patient. Loss of balance specifically after eye closure indicates impaired proprioceptive input, commonly from posterior-column or peripheral sensory disease, not a cerebellar lesion. A patient unable to stand even with visual input has ataxia, but does not thereby have this sign.

Gait is probably the most valuable noncognitive neurological assessment because it integrates motor pathways, peripheral nerves, coordination, proprioception, balance, attention, and practical confidence. Observe initiation, base, stride, arm swing, turning, symmetry, foot clearance, and response to tandem, heel, and toe walking when safe. A slow usual speed of **0.7 metres per second or less** is associated with increased risks of dementia, stroke, falls, disability, hospitalisation, and death. The number is a risk marker, not a diagnosis.

Pattern recognition then directs the differential. A broad unsteady gait suggests cerebellar ataxia. Difficulty with tandem walking may reveal mild gait ataxia. A festinating pattern suggests an extrapyramidal disorder, including a drug-induced syndrome. Circumduction of a stiff weak leg is hemiparetic. Excessive hip and knee flexion to clear the foot is a steppage pattern. Pelvic-girdle weakness produces waddling. Apparent inability to stand or walk with dramatic near-falls, yet preserved protective balance and no injury, suggests astasia-abasia. Observe gait as the patient enters the interview room.

> **GOOD TO REMEMBER**
>
> A patient encountered in bed can still have the most localising part of the examination omitted. If it is safe, watch the patient stand, start, turn, tandem, and return to the chair. The transition movements often reveal more than straight-line walking.

2.8 Psychiatric discriminators: functional signs, catatonia, and drug effects

The governing principle for a suspected functional deficit is a positive inconsistency with neuroanatomy, not an impression that the patient appears psychologically distressed. Motor and gait signs are more useful than sensory signs. Look for internal conflict between voluntary testing and automatic or synergistic movement: variable give-way effort rather than smooth yielding, a hand that avoids striking the face despite alleged paralysis, normal heel or toe walking despite marked manual ankle weakness, or a distribution that cannot be produced by a single anatomical lesion.

The **Hoover sign** tests involuntary hip extension synergy. Cup the heel of the apparently weak leg while the patient flexes the opposite hip against resistance. Downward pressure from the allegedly weak heel demonstrates preserved extension that was not produced during direct testing. The complementary pattern is absent downward pressure from the unaffected heel when the patient is asked to raise the allegedly weak leg. A related abductor manoeuvre shows reflex abduction of the “weak” leg while the unaffected leg is actively abducted. These are positive motor signs; they should be demonstrated respectfully and interpreted with the whole examination.

Functional sensory patterns include a sharp midline boundary, loss of vibration over only one half of a continuous midline bone, and pain loss with preserved temperature. Genuine proprioceptive loss should impair balance in darkness, stance with eyes closed, and relevant daily actions. Nevertheless, old follow-up series found a later neurological diagnosis in as many as **fifteen per cent** of patients initially labelled with hysteria or conversion, although those cohorts often had rudimentary examinations and investigations. Continue follow-up when the course changes.

Primitive grasping is elicited by placing the index and middle fingers in the palm. Automatic gripping, especially if unilateral, is a frontal release sign; in an inert mute patient it argues against an exclusively depressive explanation. The anatomy belongs to the Functional Neuroanatomy and Neurophysiology notes. Do not confuse this with forced handshaking in catatonia after an instruction not to take the offered hand.

Catatonia must be examined, not merely watched. Observe mutism, staring, stupor, posturing, stereotypy, echo phenomena, excitement, and withdrawal; elicit tone phenomena, negativism, waxy flexibility, automatic obedience, forced grasping, and ambitendency. Waxy flexibility begins with resistance, gradually yields to positioning, and maintains the posture. In mitgehen the limb follows the slightest pressure despite an instruction to resist; mitmachen requires more manipulation. Neurological examination is essential because catatonia can resemble drug effects, akinetic states, delirium, seizures, and advanced cognitive disorders.

TRAP **Do not treat "Gegenhalten" as a universally settled label.** A clinical tradition uses it for active resistance in advanced dementia, another for catatonic negativism in which resistance matches the examiner's force, and another associates variable resistance with frontal release. Describe what the limb did, how resistance changed with force and speed, and the accompanying signs.

In a patient receiving antipsychotic medication, examine spontaneous movement, facial animation, resting tremor, bradykinesia, rigidity, subjective restlessness with observable movement, and involuntary oral, lingual, facial, truncal, or limb movements. The Simpson-Angus Scale, Barnes Akathisia Rating Scale, and Abnormal Involuntary Movement Scale can support serial documentation, but their administration and scoring belong to the Schizophrenia: management, antipsychotics and adverse effects notes. Chorea may initially look like ordinary fidgetiness; tardive dyskinesia often involves the face, lips, jaw, and tongue; asterixis is a brief loss of sustained tone. Describe the movement before attributing it to illness, anxiety, or medication.

Examine with sensitivity. An attempt to "catch" inconsistency can worsen distress and obscure useful signs. Functional and structural disease can coexist. Emotional unconcern is not diagnostic: **la belle indifference** has no credible discriminating value and occurs with neurological disease. Base the conclusion on positive signs, anatomical coherence, longitudinal review, and appropriate investigation, not personality or bedside theatre.

Mental status validates the examination; pattern localises it; gait integrates it; anatomical inconsistency must be demonstrated, never presumed.

Lesion localisation and the imaging decision

Why it matters clinically. A psychiatric formulation can become unsafe when an observed change in language, behaviour, movement, attention, or gait is treated as a psychiatric meaning before it has first been treated as a neurologic sign. **Localisation**, also called topographic diagnosis, asks what part of the nervous system has been affected, using signs and, most often, symptoms. It is not a decorative neurological add-on to the mental-state examination. It is the disciplined step that decides whether a putative psychiatric presentation has a coherent neural address, whether it instead points away from a lesion, and what question an investigation should be asked to answer.

The psychiatrist need not reproduce a neurologist's full examination to use this reasoning. The required competence is to **appreciate neurologic signs and appraise the neurologist's conclusion**. That means separating the observation from its interpretation, resisting a psychiatric default, and recognising when a bedside pattern has more explanatory force than a familiar diagnostic label. The functional anatomy that makes a sign intelligible belongs in the Functional Neuroanatomy and Neurophysiology notes; the present task is to use that map without redrawing it.

3.1 Localisation is a step in diagnostic reasoning

Neurologic diagnosis proceeds through **four** linked operations: recognise impaired function, localise it, construct an aetiological differential, and then use ancillary procedures to decide which proposed cause is present. **Localisation** is therefore neither a substitute for diagnosis nor an inference made after a scan has supplied one. It is the bridge between a finding and a focused differential. Confusing these operations produces a familiar error: treating a radiological abnormality as the diagnosis before asking whether it accounts for the patient's observed deficit.

Recognition comes first because an unrecognised sign cannot be localised. It depends on knowing what normal performance looks like and on examining carefully enough to see a departure from it. Mild chorea may be dismissed as ordinary fidgetiness; slow eye movements may be missed if the examiner notices only whether the eyes complete the movement. Both errors begin before the localising exercise has even started. In psychiatric practice, where restlessness, unusual motor behaviour, reduced engagement, and disordered speech are often interpreted through psychopathology, this preliminary vigilance is particularly important.

Before attempting a locus, ask **three** questions: whether the reported dysfunction is real, whether it is pathological enough to warrant formal work-up, and whether it is neurologic rather than a disorder of the organ mediating that function. These are not sceptical questions directed at the patient. They are tests of the clinical observation. A sensory complaint elicited during examination is part of the examination, not mere history, and may be as objective as an observed wrist drop. Equally, a history should guide attention without hardening into preconception: it suggests whether a central or peripheral problem may be present, while the routine examination still checks all major areas.

SIGN	SITE	CAUSE	TEST
RECOGNISE IMPAIRED FUNCTION	LOCALISE THE PATTERN	FORM THE DIFFERENTIAL	CHOOSE AN ANCILLARY PROCEDURE

RULE **Give a sign its neurologic status before giving it psychiatric meaning.** Unless clear psychiatric manifestations accompany a neurologic symptom or sign, it should be taken at face value while its pattern is assessed.

This is a corrective to a durable diagnostic bias. Some movement disorders, including dystonias, were formerly cast as psychogenic before their organicity was generally recognised. The appropriate stance is neither indiscriminate testing nor reflexive disbelief. It is an accountable formulation in which a sign is first allowed to constrain the possible anatomy.

3.2 Reading the map: address and intersection

Neuroanatomy is the road map for localisation. The map has morphological structures and their functional representations: damage to a structure alters the function it mediates. Applied localisation has **two** basic routes. The **address route** begins with a characteristic constellation that directly identifies a site. The **intersection route** begins with separate signs, traces the pathway implicated by each, and locates the lesion where those pathways meet. The aim is not to recite anatomy but to make every observed finding do localising work.

REASONING ROUTE	QUESTION AT THE BEDSIDE	WHAT MAKES IT PERSUASIVE	CLINICAL CONSEQUENCE
Address	Does the whole pattern name a site?	A syndromic constellation supplies an anatomic address.	Use the pattern to narrow the differential.
Intersection	Which pathways are implicated separately?	The affected pathways converge at a shared locus.	Seek the accompanying sign that makes the location precise.
Temporal check	Were the signs acquired together?	Only contemporaneous findings may be forced into one lesion.	Separate accumulated deficits before localising.

An address-route example is the combination of resting tremor, bradykinesia, and rigidity, which points to injury of the substantia nigra. More often, a single sign only yields a range of possible locations. A third-nerve palsy, by itself, can arise along an extended course. Add contralateral hemiparesis occurring at the same time, however, and the intersection localises the process to the crus cerebri. The lesson is more general than the particular syndrome: one sign opens a corridor; the added, anatomically coherent sign identifies the room.

MNEMONIC

PATTERN - PATHWAY - PLACE - PHASE

Pattern asks whether the signs already provide an address. If not, identify the **pathway** indicated by each finding, locate their shared **place**, and check the **phase** of onset so old and new deficits are not falsely intersected. This is a reasoning scaffold, not a replacement for anatomy.

The temporal check prevents an elegant but false conclusion. The examination can display the accumulated end result of multiple lesions. Signs acquired in widely separated periods should not be made to intersect at a single locus simply because they coexist today. Here the history has a restricted but decisive localising role: it determines whether findings belong to the same event. Otherwise, the history chiefly contributes aetiology, whereas the examination carries the greater localising weight.

3.3 What the examination localises, and what history explains

The examination has greater localising value than the history. By tracking back the pathways mediating impaired functions found at the bedside, a lesion can often be localised even when the history is sparse or unreliable. This is especially useful in liaison work, where collateral may be incomplete, the patient may be distressed or unable to cooperate, and the narrative may be shaped by the very disorder that needs explanation. It does not diminish the history. It assigns it its proper job.

The temporal evolution of an already localised deficit helps define cause. The same left-sided hemiparesis carries different aetiological implications when it develops over minutes, over days, or insidiously over months: cerebrovascular disease or epilepsy, infection or demyelinating disease, and tumour or a degenerative process respectively. The anatomical conclusion remains the same; the causal differential changes. This division of labour stops the clinician from allowing a dramatic tempo to dictate an anatomical site that the examination has not supported.

- **History:** establish onset, evolution, and whether apparently related signs belong to one clinical event.
- **Examination:** identify the impaired functions and construct the anatomical pattern.
- **Formulation:** combine locus and tempo into a restrained differential that includes plausible disease and dangerous alternatives.

The motor examination illustrates the distinction. Severity tells the clinician about functional capacity, but the pattern of weakness supplies more localising information. The level-by-level pattern itself is addressed in the Functional Neuroanatomy and Neurophysiology notes. At this stage, the psychiatrist should read a documented motor pattern as evidence about site, not merely as a measure of disability.

A useful formulation also travels along **two axes**: the affected level, broadly central nervous system, peripheral nervous system, or muscle; and the distribution, focal versus diffuse. Together, site and extent direct the differential. Discrete lesions are the classical terrain of localisation, with cerebrovascular disease the commonest cause and demyelination, infection, trauma, and tumour also often behaving as discrete lesions. System lesions instead affect arrays of neurons serving a functional system. The method remains useful for them, but the symptom pattern may point to an affected pathway rather than the primary cellular target. This is why a localising conclusion should be stated as an inference from signs, not as an overconfident claim that the disease process has been anatomically exhausted.

3.4 Higher functions, posture, and gait demand pattern recognition

Precision generally decreases as one ascends the neuraxis. Lower-level lesions tend to produce stable findings, whereas higher-level lesions may vary even within a single examination. The correct question for a higher function is therefore not simply whether the patient can perform it at some moment, but whether the patient performs it consistently in a normal way. Higher networks can mask a deficit, so sampling must be sufficient to defeat that masking. This is a clinical discipline, not a demand for a scored cognitive screen; formal screens and their interpretation belong in the Psychometry, Assessment and Descriptive Psychopathology notes.

In practice, a patient who manages a task once but fails it unpredictably has not thereby demonstrated intact function. For native speakers, any difficulty repeating a sentence should be regarded as abnormal. Record the performance itself, its consistency, the patient's engagement, and the accompanying neurologic signs. Such documentation helps the next clinician distinguish fluctuating capacity, language impairment, impaired attention, and a failure of testing conditions without pretending that a single bedside answer settles the matter.

Posture and gait need a related but distinct approach because they are assessed chiefly by functional observation, rather than by isometric strength testing or sensory questioning. Their localising value comes from characterising the movement or postural pattern and seeking accompanying neurologic signs. Some patterns are highly stereotyped and strongly localising; cautious gait and central disequilibrium, by contrast, have many potential causes and are harder to localise. The examination technique for gait and the rest of the physical neurological examination are addressed elsewhere in this chapter; here the point is to use the observed pattern as a hypothesis, then test that hypothesis with associated signs.

PITFALL

Do not call a gait or postural complaint non-neurologic merely because muscles and joints can also alter movement. Clinical bias tends to favour a non-neurologic diagnosis when the neural control of posture or gait is actually impaired. Observation must be followed by a search for a coherent accompanying neurologic pattern.

3.5 Formulation includes the limits of localisation

The formulation should answer the **Four Questions of Neurology**: what are the symptoms, what are the signs, where is the lesion, and what is the lesion. In psychiatric work, the equally important negative conclusion is that neurologic disease does not explain the presentation. Neither conclusion is reached by impression alone. The clinician must show how the signs cohere anatomically, or why they fail to do so, and must keep the resulting differential proportionate to the evidence.

Localisation is valuable in most cases, but it is partly an art and is not applicable to every important neurologic illness. Alzheimer disease is an example of a setting in which the classical focal-lesion method is not the right intellectual frame. Similarly, widespread or symmetric changes should not be forced into a discrete lesion formulation. The method is strongest when it clarifies a specific pattern, less decisive where disease is diffuse, degenerative, or network-based. A

good examiner acknowledges this boundary rather than manufacturing a misleading anatomic precision.

The differential begins with diseases most consistent with the signs and locus, includes unlikely but potentially life-threatening conditions, and avoids the intellectual flourish of unlikely but fascinating explanations. This is not anti-curiosity. It is an insistence that the differential stay connected to the patient rather than to the examiner's imagination. A concise written formulation makes that connection visible and gives imaging or other testing a defensible purpose.

TRAP **A single abnormality is not a licence to announce a precise lesion.** One sign often underdetermines the locus. Seek a second anatomically coherent sign, then ask whether the findings appeared in the same clinical period before applying the intersection route.

Conversely, an anatomically impossible deficit points away from a lesion locus. Preserved temperature sensation with loss of pain perception is nonanatomic and should prompt a careful consideration of psychogenic origin. This is the inverse of localisation, not a shortcut to disbelief. Long follow-up of classic hysteria and conversion cohorts found specific neurologic conditions emerging in **up to 15%** of patients. Those cohorts had limited examinations and testing, but the practical warning remains: a psychogenic conclusion must be positively grounded in the pattern, never merely inferred from psychiatric context or diagnostic inconvenience.

3.6 The liaison problem: language that looks psychiatric

The psychiatric presentation may be the sign. Sudden fluent aphasia can account for apparent anger, acute psychosis, consternation, behavioural change, or confrontations around management in liaison settings. Its disordered output may also resemble psychotic speech. A clinician who sees bizarre language and assumes thought disorder can miss the neurologic basis of the presentation, particularly when there is no obvious hemiparesis and attention cannot be captured well enough to examine sensation or visual fields.

The discrimination is clinical and longitudinal. Schizophrenic speech usually develops gradually in the context of long-standing psychiatric illness; language abnormalities can become increasingly similar to aphasia as thought disorder develops. Fluent aphasia usually presents suddenly, may occur with stroke risk factors, and often leaves the patient aware that communication is failing, seeking help, and giving short responses. Repetition of polysyllabic words and complex phrases is an important discriminator because people with schizophrenia can do this unlike most people with fluent aphasia. None of these observations should be isolated into a mechanical rule. Frightened, confused people with sudden aphasia may themselves be agitated or irrational.

QUESTION	FLUENT APHASIA PATTERN	PSYCHOTIC-SPEECH PATTERN	WHY IT MATTERS
Tempo	Usually sudden	Usually gradual	Tempo changes the aetiological frame.
Repetition	Often impaired	May be retained	Test language, not only content.
Patient stance	May recognise communication failure and seek help	Interpret in the broader psychiatric history	Behaviour alone is not diagnostic.
Associated signs	Seek reflex, sensory, and visual-pathway clues	Do not assume their absence from a brief encounter	Associated signs can localise the problem.

The localising endpoint is usually a small discrete structural lesion in the temporoparietal region, often a stroke, although neurodegenerative illnesses can also cause fluent aphasia. The detailed focal lobe syndromes belong in the Stroke, TBI, delirium and focal brain syndromes notes. The immediate psychiatrist's task is simpler and more urgent: recognise that the speech may be language failure, document the pattern, look for accompanying signs, and frame the investigation around that localisation.

GOOD TO REMEMBER

"Gibberish" is a description, not a diagnosis. When language is suddenly abnormal, test repetition and seek a neurologic pattern before attributing the communication failure to psychosis, personality, or non-cooperation.

3.7 From a bedside locus to the imaging question

Ancillary procedures should follow clinical diagnosis if they are to realise their diagnostic potential. This is not an argument against imaging, neurophysiology, or molecular investigation. It is an argument against using them as an unstructured search. **The locus reached on examination names the region that the scan must interrogate**. A patient with a left-sided hemiparesis and a right third-nerve palsy that appeared together, for example, requires careful attention to the midbrain when MRI is obtained. The examination has put a question to the study.

That ordering protects against false-positive attribution. A broad, unlocalised approach to neuroimaging may reveal a striking congenital brain cyst and wrongly make it responsible for a neurologic disorder while the actual disease remains overlooked and untreated. Careful localisation is the guard against this error. It also makes ancillary procedures more thoughtful, reducing patient discomfort and wasted resources. The question is not, "What can a scan find?" but, "Does the study answer the anatomical and aetiological question posed by this patient's signs?"

Some diagnoses cannot be closed at the bedside alone. Stroke and brain tumour require imaging studies, and confirmation of seizures often requires EEG. The modality, its indications,

sequences, and red-flag lens are covered in the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes and should not be reconstructed from this section. Here, the decision workflow is sufficient: establish the clinical syndrome, state the likely locus and differential, and request the study that can adjudicate that framed question.

Localisation also determines what kind of diagnostic question is being asked. Aphasia usually implies a discrete dominant cerebral-hemisphere injury and therefore raises a structural question for CT or MRI. Dementia instead raises a neurodegenerative question, with blood tests as well as CT or MRI in its evaluation. The distinction is not academic: it changes the evaluation before it changes the label. When an imaging study shows a different process from the clinical formulation, abnormal scan findings routinely override the clinical impression. This should encourage humility, not passivity: a clear bedside formulation makes a discrepant result intelligible and helps the clinician revisit both the patient and the image.

> Localise from the bedside pattern first; then make every investigation answer the question that localisation has already defined.

4 · Electroencephalography in psychiatry: normal rhythms, common abnormalities and clinical indications

Why EEG still matters to the psychiatrist. Electroencephalography is not a laboratory badge for a psychiatric diagnosis. Its value lies in a disciplined change of question: from "Which psychiatric disorder is this?" to "Could an epileptic, encephalopathic, toxic, or other cerebral process be contributing to this presentation?" That distinction protects patients from both over-investigation and premature psychologising. The normal frequency bands, the morphology of epileptiform discharges, and the distinction between interictal and ictal phenomena are established in *the Epilepsy and psychiatry notes*; the present discussion applies that vocabulary to the psychiatric workup.

EEG is a real-time sample of cortical electrical activity, not a general screen for brain disease. A useful request therefore begins with a clinical problem, a timeline, medication exposure, and the precise alternative diagnosis being tested. In uncomplicated psychiatric illness, recordings are commonly **normal or only nonspecifically abnormal**, with such findings as excess beta or theta activity, isolated sharp waves or spikes, or mild disorganisation. These patterns do not convert a clinical syndrome into a neurological diagnosis.

4.1 Recording language and the meaning of a normal study

Before interpreting an abnormality, the psychiatrist should be able to read the report's spatial language. The **10-20 International System** places electrodes according to measured proportions of head landmarks, so that comparable labels refer to comparable cortical regions despite differences in head size. Electrodes are positioned at **10% or 20%** of the relevant measured distances. This is why serial recordings and reports from different laboratories can be compared meaningfully rather than treated as idiosyncratic maps.

The labels are a clinical shorthand, not a claim that an electrode records only one circumscribed cortical site. In the standard nomenclature, F, T, C, P and O denote frontal, temporal, central, parietal and occipital locations; *z* identifies the midline; odd numbers are left sided and even numbers right sided. Cz is the vertex. The conventional system comprises **21** electrodes, including scalp and aural reference sites. A psychiatrist need not become an electroencephalographer, but should recognise when a report is describing a focal distribution, a generalised background change, or a technically interpretable routine recording.

GOOD TO REMEMBER

Electrode names describe a standardised scalp location. They help communicate distribution; they do not by themselves establish the aetiology of a psychiatric presentation.

A normal EEG is similarly contextual. It can be highly informative when the clinical question is diffuse toxic-metabolic encephalopathy, yet it cannot reliably detect or exclude a structural lesion. When structural pathology is being considered, neuroimaging is the standard investigation; a normal EEG must not be used as a structural rule-out. The relevant imaging approach is covered in *the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes*.

4.2 Indication-driven use in general psychiatry

The governing principle is that a **routine EEG screen is unwarranted** in psychiatric practice. There is no accepted diagnostic EEG indication for psychiatric disorders as defined in the current diagnostic manuals. The test earns its place when it alters the differential towards an organic cause or answers a sharply framed seizure-related question. Asking for it simply because symptoms are severe, unusual in content, or treatment-resistant without further reasoning invites nonspecific findings and false certainty.

RULE **Order EEG to test an organic alternative, not to certify a psychiatric diagnosis.** The clinically useful psychiatric role is differential diagnosis of organic aetiology and investigation of sleep disorders, rather than diagnostic confirmation of a DSM syndrome.

Its yield rises when there is an **atypical presentation, atypical age at symptom onset, or atypical treatment response**. In that setting, the report is sought to identify paroxysmal activity or focal or diffuse slowing that would redirect the assessment. This does not mean every late presentation or difficult-to-treat illness needs a recording. It means the clinician should articulate why the phenomenology, course, or treatment response makes a cerebral contribution plausible and what the result would change.

CLINICAL QUESTION	EEG CONTRIBUTION	INTERPRETIVE BOUNDARY
Uncomplicated psychiatric syndrome	Usually none beyond occasional nonspecific findings	Do not use as a routine screen
Atypical course or response	Seek paroxysmal activity or focal/diffuse slowing	Interpret alongside the syndrome and examination
Acute agitation or unresponsiveness	Detect diffuse slowing or ongoing epileptic activity	Acts on the immediate organic differential
Cognitive decline with selected warning features	Supports a specific encephalopathic or periodic pattern	Not a routine dementia test

A **paroxysmal EEG finding** alone does not diagnose epilepsy. The retrospective clinical history must still support an epileptic nature for the events. In an atypical psychiatric presentation, such a finding should lower confidence in a purely psychiatric formulation and prompt re-examination of the history, collateral account, medication list, and neurological assessment. It is a reason to think again, not a licence to relabel the patient.

TRAP **"Spikes equal epilepsy, therefore the presentation is not psychiatric."** EEG and clinical semiology must be reconciled. An isolated paroxysmal finding changes diagnostic confidence; it is not independently diagnostic.

4.3 Acute disturbance: delirium, toxicity and nonconvulsive seizures

In an acutely agitated, fluctuating, or nonresponsive patient, EEG can be immediately decisive because the key distinction is between a delirious encephalopathic state and ongoing epileptic activity. **Diffuse slowing** supports delirium, while evidence of **nonconvulsive status epilepticus** demands a different urgent pathway. The accompanying figure places these patterns in the broader psychiatric workup. This is an investigation of the acute brain state, not an adjunct to a routine mental-state examination.

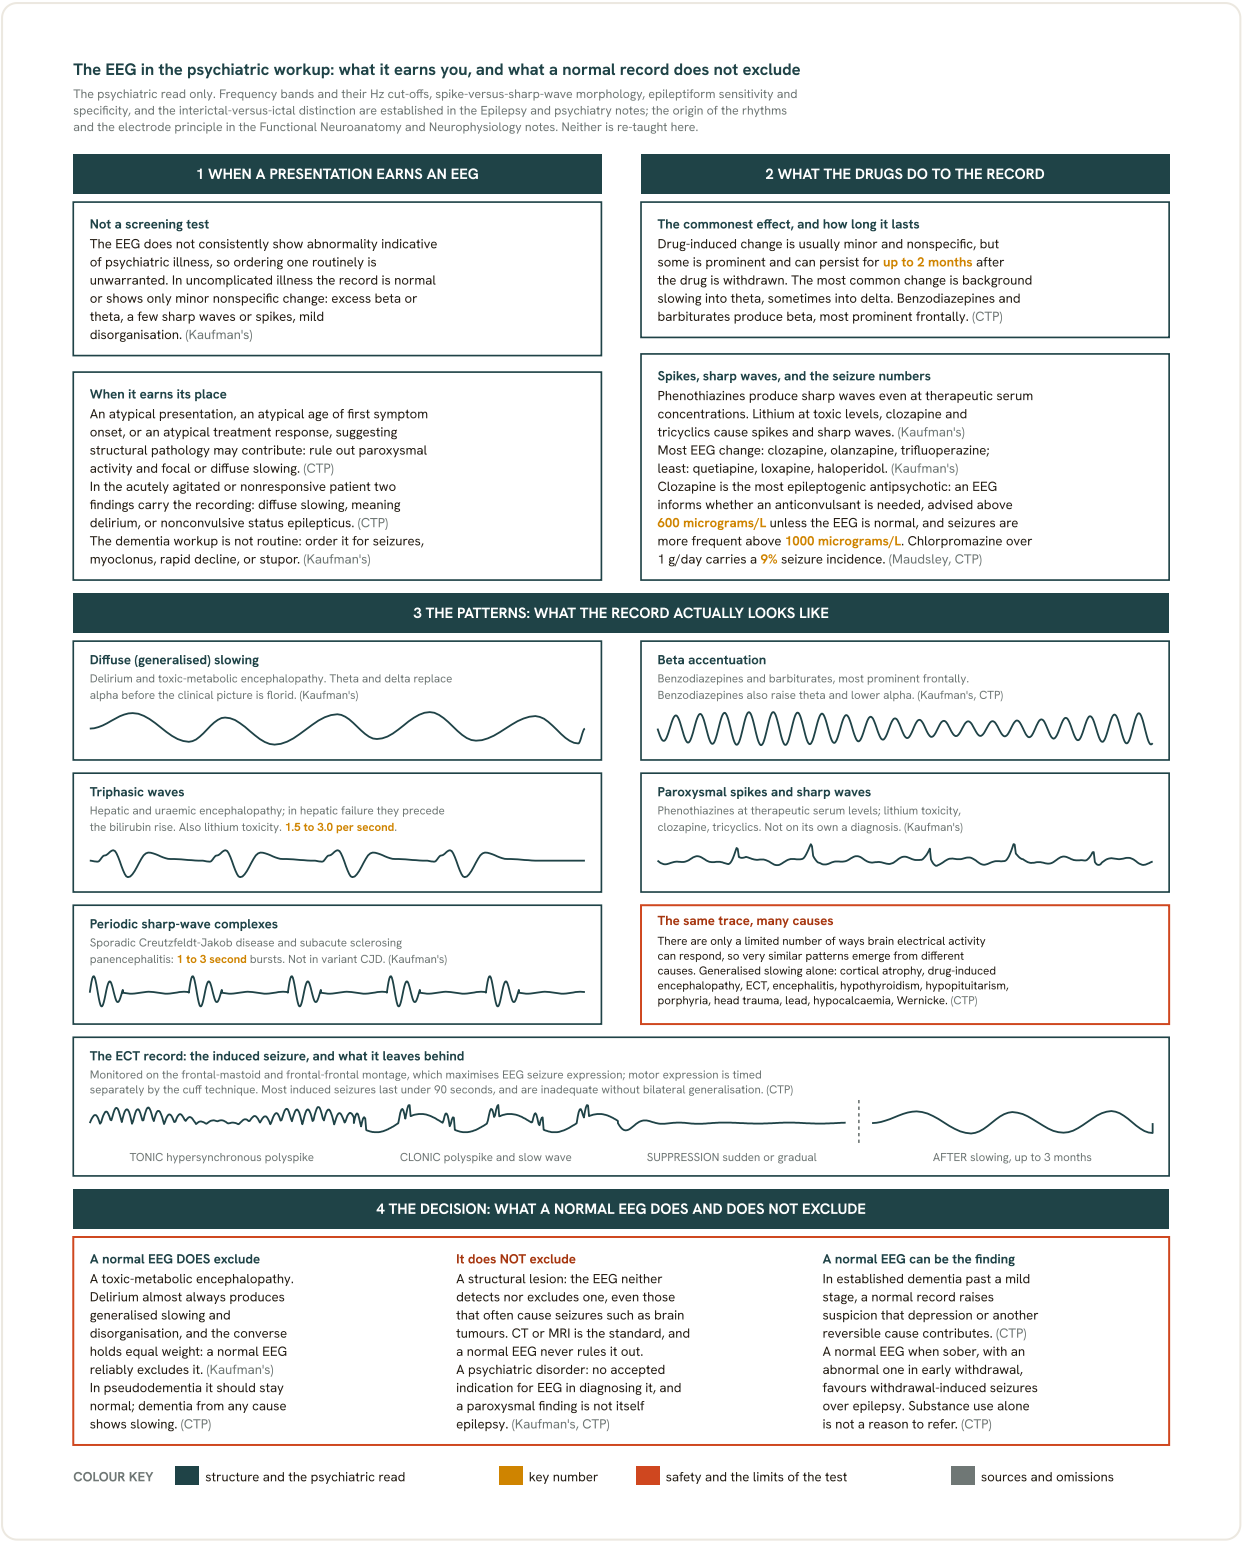


Fig 31.2 The EEG in the psychiatric workup: diffuse slowing, triphasic waves, and the psychotropic and ECT record. The decision is the teaching, not the tracing. The EEG is not a screening test and has no accepted role in diagnosing a psychiatric disorder; it earns its place when the presentation, age of onset or treatment response is atypical, and when an agitated or nonresponsive patient may be delirious or in nonconvulsive status epilepticus. Read the patterns for what they exclude: a normal record reliably excludes a toxic-metabolic encephalopathy, but never a structural lesion and never a psychiatric diagnosis. Triphasic waves are the one shape pointing somewhere specific, and lithium toxicity is their mimic. EEG abnormality in up to 20% of healthy individuals appears nowhere here as fact: its source calls it a belief still needing study. The EEG fundamentals belong to the Epilepsy and psychiatry notes, and the origin of the rhythms and the electrode principle to the Functional Neuroanatomy and Neurophysiology notes. (Kaufman's Clinical Neurology for Psychiatrists; Kaplan and Sadock's CTP; The Maudsley Prescribing Guidelines.)

In toxic-metabolic encephalopathy, theta and delta activity can replace alpha early, sometimes while behavioural or cognitive change is still subtle; the typical record shows generalised slowing and disorganisation. Background slowing is nonspecific because many different processes converge on the same electrical response. It should therefore trigger a targeted clinical search rather than a satisfied conclusion. A normal EEG, however, **reliably excludes toxic-metabolic encephalopathy**, making the test especially useful when delirium is genuinely in the differential.

Triphasic waves are a more suggestive pattern: high-voltage, frontally dominant or diffuse slow waves with a characteristic blunt or rounded spike-like onset, occurring at **1.5 to 3.0 Hz**. They often point toward hepatic or uraemic encephalopathy and, in hepatic failure, may precede the bilirubin rise. The full pattern may also include frontal triphasic activity with loss of organised occipital alpha. Yet the psychiatrist must retain the medication differential: lithium toxicity is a recognised cause of triphasic waves, and metabolic causes remain common.

PITFALL

Do not dismiss diffuse slowing because a psychotropic plasma concentration is within the therapeutic range. Drug-induced encephalopathy can show diffuse slowing despite a therapeutic serum level, so toxicity remains a clinical possibility alongside metabolic and endocrine causes.

Psychotropics complicate this reading. Significant diffuse slowing in an unstable patient receiving psychotropic medication, particularly an older adult, should prompt consideration of medication toxicity as well as alternative causes of encephalopathy. Lithium at intoxicating levels can be associated with diffuse slowing, paroxysmal spikes and triphasic waves, and these EEG abnormalities may remain after the serum concentration has normalised. The EEG is therefore a measure of cerebral effect, not merely a surrogate for the current blood value.

21	1.5–3.0 Hz	LESS THAN 90 s
STANDARD ELECTRODES	TRIPHASIC WAVES	MOST ECT SEIZURES

4.4 Medication effects: read the drug exposure before the tracing

Medication history is part of EEG interpretation, not a footnote. Psychotropic-induced changes are usually minor and nonspecific, but some persist for up to 2 months after withdrawal. The commonest effect is background slowing into theta, sometimes delta. A report should therefore be read with the current regimen, recent discontinuations, dose changes, intoxication risk, and clinical state in view. Generalised slowing is not a diagnosis, and its differential includes drug-induced encephalopathy, ECT and systemic illness.

EXPOSURE	CHARACTERISTIC EEG TENDENCY	CLINICAL READING
Benzodiazepines	Beta enhancement, most prominent frontally; slight theta increase and alpha decrease	Can reduce epileptiform activity
Barbiturates	Frontal beta during use; paroxysmal activity and spikes after abrupt withdrawal	Withdrawal discharges may be behaviourally silent
Phenothiazines	Sharp waves	May occur at therapeutic concentrations
Lithium toxicity	Diffuse slowing, spikes and triphasic waves	Abnormality can outlast normalisation of the serum level
Typical antipsychotics	More alpha, lower frequencies and less beta	Spikes and seizures may occur at high doses
Tricyclic antidepressants	Slow and fast rhythm increase, lower alpha frequency, paroxysmal slow waves and spikes	Seizures may occur at higher doses
SSRIs	Less alpha and more theta	Lower seizure rates than tricyclics

The table is a pattern-recognition aid, not a prescribing hierarchy. Benzodiazepine-related **frontal beta activity** is a classic sedative-hypnotic signature. Barbiturate withdrawal can produce generalised paroxysmal activity and spikes for up to 3 to 4 weeks after the last exposure, sometimes without overt motor seizures. Medication effects can thus mimic disease, expose withdrawal physiology, or coexist with the illness that led to treatment.

The operational consequence is modest but important. An abnormal result should prompt a review of temporal relationships: did the change appear after introduction, dose escalation, toxicity, abrupt discontinuation, or ECT? Does it fit the patient's vigilance, cognition and neurological signs? A plausible drug effect may explain a trace without explaining the whole patient. Conversely, attributing slowing to a familiar prescription can obscure concurrent systemic illness or nonconvulsive seizure activity. The report becomes useful when it is treated as evidence within a formulation rather than as a pharmacological fingerprint printed in isolation.

Clozapine requires particularly careful integration of clinical, EEG and plasma-level information because it is the most epileptogenic antipsychotic and EEG abnormalities are common in treated patients. It can increase delta, theta and beta activity over frontal, central and parietal areas, with paroxysmal slowing and spikes. An EEG may help determine whether anticonvulsant treatment is needed, although the interpretation is complex. At plasma levels above **600 micrograms/L**, anticonvulsant use is recommended unless the EEG is normal; seizures occur more frequently above **1000 micrograms/L**. This is a deliberately scoped clozapine-management decision, not a general threshold for abnormal EEGs.

MNEMONIC

Drugs before discharges

Before treating an EEG discharge as disease, ask what the patient has received, stopped, accumulated, or cleared. The medication list often explains the trace, but it must never replace assessment for concurrent encephalopathy or seizures.

4.5 EEG during and after electroconvulsive therapy

EEG monitoring during ECT documents the induced seizure and helps assess its cerebral expression. The ictal EEG resembles a generalised tonic-clonic seizure and its aftermath. Frontal-mastoid and frontal-frontal montages are used to maximise seizure expression; during the tonic phase, hypersynchronous polyspike complexes are followed by polyspike-and-slow-wave complexes, and the clonic phase is followed by suppression. Motor expression is assessed separately, so an EEG record is not merely a more technical version of observing movement.

Most induced seizures last less than 90 seconds, but duration alone does not establish adequacy. Bilateral generalisation on EEG is the stated adequacy criterion. This distinction matters when a seizure appears brief, or when motor expression is limited by anaesthesia and muscle relaxation: the clinical team needs evidence of a generalised cerebral seizure, not a stopwatch result isolated from the tracing.

After ECT, slow-wave activity can develop over frontal regions or the whole cerebrum and persist for up to 3 months. With unilateral treatment, slowing may remain restricted to the treated side. Postictal suppression and slow-wave activity are more likely in responders than nonresponders. Slow-wave change is not simply an adverse marker: it is associated with both memory impairment and effective treatment of depression. Interpretation should therefore acknowledge a benefit-cost signal rather than treating post-ECT slowing as automatically pathological.

4.6 Cognitive decline and selected disease signatures

EEG is not indicated routinely in dementia assessment. Common dementia-producing illnesses tend to produce slowing or other minor nonspecific abnormalities, which may be difficult to distinguish from age-related change. The recording becomes more useful when seizures, myoclonus, rapidly progressive decline, stupor, suspected nonconvulsive seizures, toxic-metabolic encephalopathy, or a dementia with a specific EEG pattern is under consideration.

In early Alzheimer disease, background activity usually slows to below 8 Hz; later records are slow and often disorganised. Vascular dementia can also cause abnormalities, but EEG cannot reliably differentiate it from Alzheimer disease. Conversely, a normal EEG in a patient thought to have dementia beyond a mild stage should raise suspicion of depression or another reversible contributor. Depression-related cognitive impairment may have a normal or only mildly slowed EEG, but EEG cannot quantify the relative contribution of depression and mild dementia when both are present. The fuller diagnostic framework belongs in *the Dementias and neurocognitive disorders notes*.

The striking exception is the periodic sharp-wave complex. In sporadic Creutzfeldt-Jakob disease and subacute sclerosing panencephalitis, dementia and myoclonus may be accompanied by **periodic sharp-wave complexes**. Classically, fairly regular bursts occur across channels at intervals of 1 to 3 seconds, with low activity between bursts; myoclonic jerks are their clinical counterpart. Variant Creutzfeldt-Jakob disease does not show this characteristic periodic pattern. The clinical context remains decisive, but this is one of the rare psychiatric-neurological differentials in which EEG can provide a near-definitive signature.

4.7 Substance use, ageing and a practical reporting discipline

Drug use alone is not a sufficient indication for EEG referral. With infrequent exceptions, recreational or dependent use does not introduce frank abnormalities into the visually analysed tracing. The question becomes useful when there are seizures, altered consciousness, a concerning withdrawal state, or an alternative neurological diagnosis. In alcohol-related seizures, a normal EEG during sobriety with an abnormal early-withdrawal EEG favours withdrawal-induced seizures rather than epilepsy with comorbid alcohol dependence. The substance-use assessment itself belongs in *the Substance use disorders notes*.

A practical report-reading discipline prevents two opposite errors. First, do not inflate a minor nonspecific change into an explanation for hallucinations, mood symptoms or personality change. Second, do not call a record “nonspecific” and stop when the patient is clinically delirious, fluctuating, or difficult to arouse. The EEG contribution is strongest when its finding is converted into a focused next question: diffuse slowing asks about global cerebral dysfunction; a periodic pattern asks about a narrow diagnostic class; apparent epileptic activity asks whether the witnessed events and collateral history are epileptic. This preserves the hierarchy in which the patient, not the tracing, remains the object of diagnosis.

Older psychiatric patients deserve a lower threshold for organic thinking when seizures are suspected. Older adults may account for as many as one-fourth of new-onset epilepsy cases. This does not make EEG a screening test for every late-life psychiatric presentation. It means that late onset, episodic phenomena, fluctuating awareness, unexplained falls, atypical behaviour, and abrupt cognitive change should be actively tested against an organic differential rather than absorbed into a familiar psychiatric label.

When reviewing or requesting an EEG, document the question in a sentence: “Is there evidence of diffuse encephalopathy, ongoing epileptic activity, or a focal abnormality that would alter the formulation?” Then document the clinical state during recording and relevant medication exposure. If the study is normal, state what it does and does not reduce in the differential. If it is abnormal, distinguish a nonspecific background change from a disease-linked pattern, and return to the patient rather than allowing the report to become the diagnosis.

A well-framed request also makes the eventual report safer to use. Give the electroencephalographer the target symptom, the temporal pattern, recent events, medication and substance exposure, and whether the patient was confused, drowsy or clinically recovered during the recording. On return, reconcile the result with the original question before changing

the psychiatric formulation. That small discipline prevents an incidental tracing feature from acquiring more authority than the history.

- State the clinical phenomenon that prompted the request and the organic alternative under consideration.
- Record the patient's state during the study and every exposure that could alter the background or provoke a withdrawal effect.

> EEG in psychiatry is an indication-driven test: use it to detect cerebral processes that change the formulation, never to validate a psychiatric label.

5 · Cerebrospinal fluid examination and its role in the neuropsychiatric workup

CSF examination is a question-driven investigation, not a routine psychiatric screen. Its value lies in sampling the biochemical and cellular environment of the central nervous system when infection, inflammation, bleeding, pressure disturbance, or a selected neurodegenerative process is plausible. The marks are therefore earned less by reciting a laboratory list than by showing sound sequence: decide whether the result can change the differential, establish that puncture is safe, do not postpone time-critical treatment, request the right studies, and interpret the pattern in its clinical context.

Lumbar puncture is the usual route to CSF. It is invasive and uncomfortable, and its commonest complication is post-lumbar puncture headache. That simple fact should discipline requesting behaviour without making the test seem exceptional or mysterious. A well-chosen study may convert a broad neuropsychiatric presentation into a treatable infectious or inflammatory diagnosis; a poorly chosen study exposes the patient to discomfort and risk while producing results that are difficult to interpret.

5.1 Start with the clinical question

The first decision is not "Which CSF tests should I order?" but "What process am I trying to confirm or exclude?" Acute headache with fever and neck stiffness raises a different question from subacute behavioural change, progressive cognitive decline, or papilloedema. CSF is most useful when the differential contains a process that directly alters pressure, cells, protein, glucose, immunoglobulins, microbial material, or disease-associated proteins. It is much less useful as a broad fishing exercise after an unrevealing psychiatric assessment.

A classic bedside trigger is the presence of at least **two of three** features: headache, fever, and nuchal rigidity. In that setting, puncture is directed toward meningitis, subarachnoid haemorrhage, or another inflammatory disorder of the central nervous system. The triad is a prompt to act, not a substitute for clinical judgement. Neuropsychiatric presentations may foreground confusion, altered behaviour, reduced engagement, or cognitive change; the psychiatrist must still look actively for neurological and systemic clues rather than waiting for a fully expressed textbook syndrome.

The **emergency CSF indication** is central nervous system infection, particularly meningitis. Encephalitis also requires urgent analysis and is managed presumptively as herpes simplex encephalitis until CSF polymerase chain reaction is negative. This distinction matters because diagnostic urgency and treatment urgency are coupled but not identical: the sample should be obtained promptly when safe, while treatment must not be held hostage to the procedure.

RULE **Send CSF only when the clinical question can be stated before the needle is inserted.** "Exclude organicity" is not a sufficient request; name the suspected compartmental process and the result that would alter management.

5.2 Safety before puncture, treatment without delay

The decisive hazard is not raised pressure in the abstract. It is a pressure gradient created by a mass lesion or obstructive hydrocephalus. Removing CSF from the spinal compartment can lower pressure below the lesion while intracranial pressure remains unopposed, precipitating transtentorial displacement. The feared outcome is **brain herniation** and death. This is why the pre-procedure neurological examination is not ceremonial paperwork: fundoscopy, arousal, and focal signs determine whether imaging must come first.

Pre-lumbar-puncture imaging is mandatory when papilloedema, any impairment of consciousness, or any focal neurological sign is present. A mass lesion or obstructive hydrocephalus contraindicates puncture. Local sepsis matters as well: a sacral decubitus ulcer contraindicates the procedure because the needle may carry organisms into the spinal canal. These are different mechanisms of harm, and each must be checked rather than compressed into the vague instruction "rule out contraindications."

MNEMONIC

Fundi, Consciousness, Focality, Compartment, Skin

Before puncture, check the **fundi** for papilloedema, **consciousness** for any impairment, the examination for **focality**, imaging for a dangerous intracranial **compartment** lesion when triggered, and the puncture-site **skin** for infection.

Sequence becomes especially important in suspected pyogenic bacterial meningitis. If clinical assessment finds no contraindication, puncture follows expeditiously after assessment and blood cultures. If imaging is required, antibiotics begin after assessment and blood cultures rather than waiting for the scan or the subsequent puncture. CSF obtained after imaging has established safety may remain diagnostically useful even after treatment has begun.

PITFALL

Do not turn "image before puncture" into "image before antibiotics." In suspected bacterial meningitis, treatment must not wait for imaging or lumbar puncture. The safety sequence protects the patient from herniation; it does not justify delaying antimicrobial treatment.

Outside suspected acute bacterial meningitis or subarachnoid haemorrhage, the usual sequence is to postpone puncture until imaging has excluded an intracranial lesion when one is a concern. In those rapidly evolving conditions, diagnosis is time-sensitive, but the bedside safety screen still governs whether puncture can proceed. “Urgent” never means “blind.”

5.3 What the procedure is meant to capture

The needle is inserted caudal to the lower boundary of the spinal cord to avoid cord injury. The procedural objective is not merely to obtain clear fluid. Depending on the question, the study may need a pressure measurement, visual inspection, cell counts and differential, chemical analysis, microbiological studies, or specialised protein and immunological tests. The request form should therefore carry the working diagnosis and the clinically relevant timing, rather than the unhelpful instruction “CSF routine.”

Opening pressure is essential when papilloedema persists despite a good-quality magnetic resonance study that does not show a structural cause. In that context, pressure measurement and CSF analysis are required together: the former pursues idiopathic intracranial hypertension, while the latter helps exclude low-grade chronic meningitis. Raised pressure without a mass lesion is not itself the same contraindication as a dangerous pressure gradient; after appropriate imaging, puncture in idiopathic intracranial hypertension is diagnostic and may occasionally be therapeutic.

The **routine CSF profile** comprises colour, red and white blood cell counts, protein, and glucose. These measurements answer different questions. Appearance may disclose blood or pigment; the cell count establishes whether there is pleocytosis and the differential suggests the inflammatory response; protein reflects disruption and inflammatory burden; glucose must be interpreted with blood glucose. The accompanying figure places these core constituents alongside opening pressure and oligoclonal immunoglobulins, and links them to the decision to send CSF.

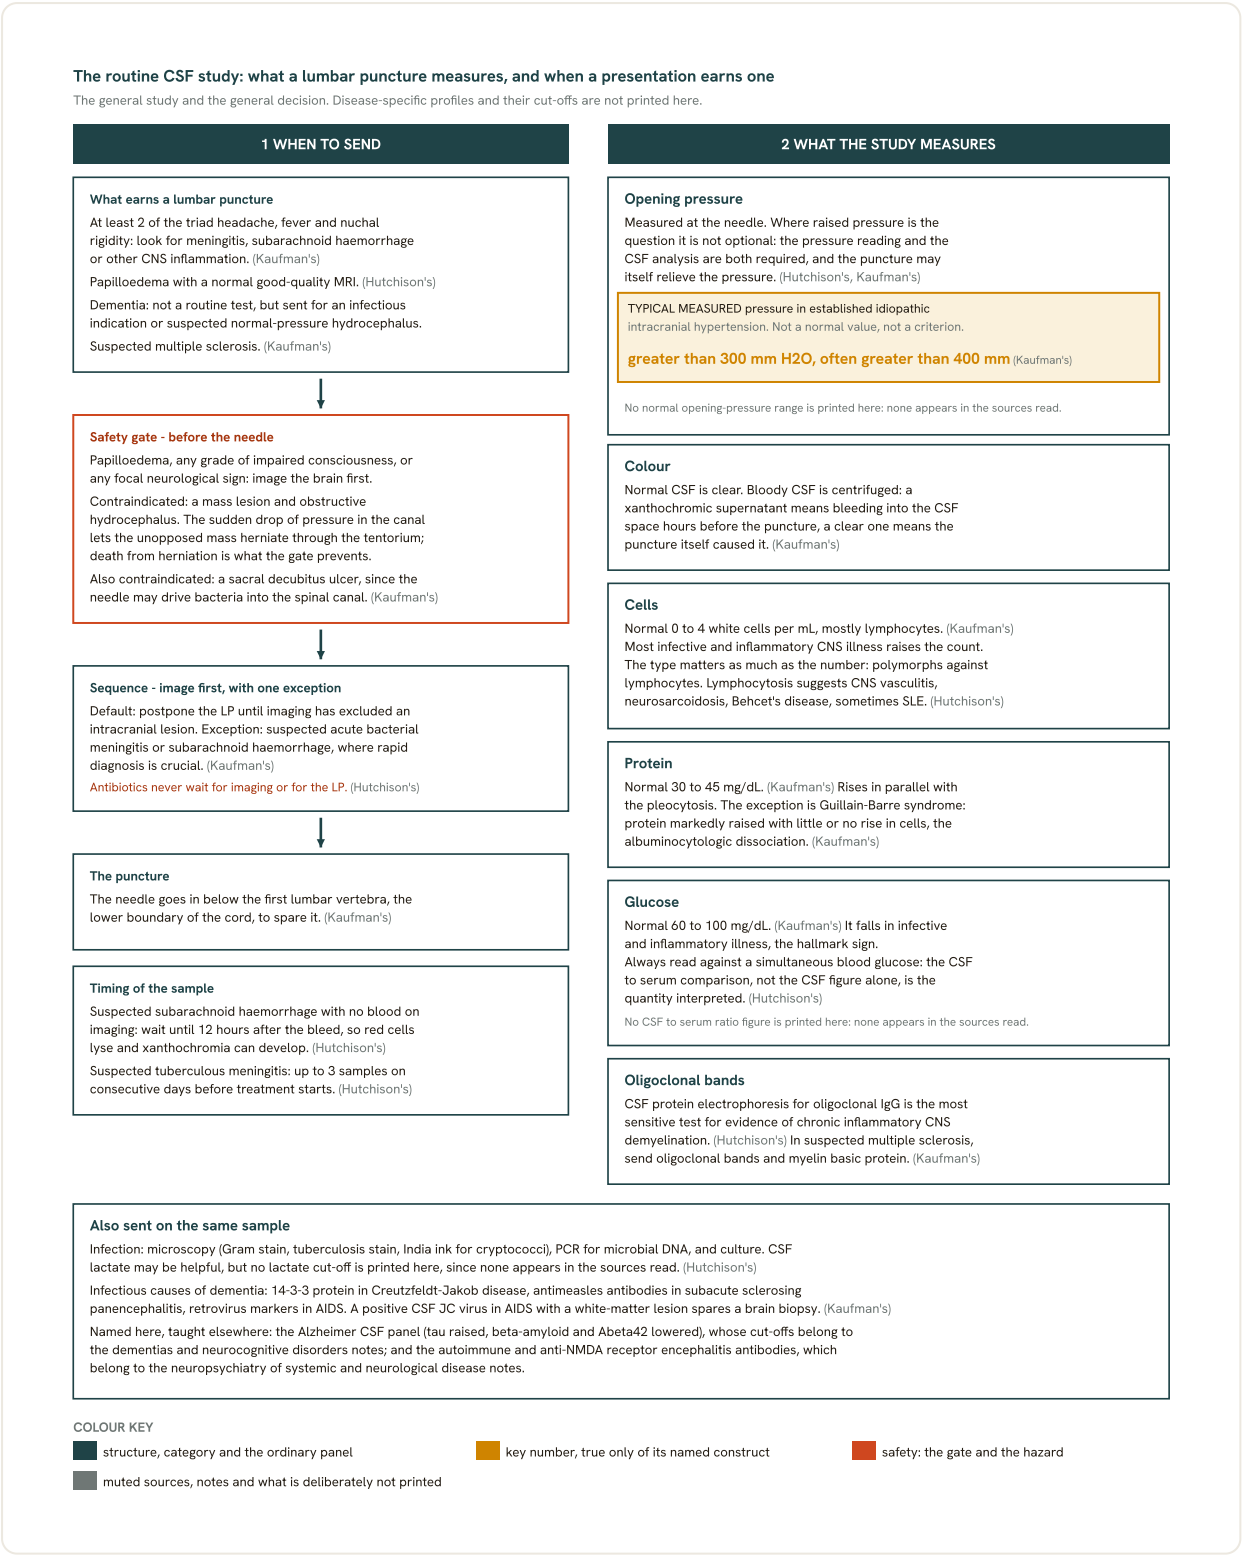


Fig 31.3 The routine CSF study: the constituents measured and the general when-to-send logic. The decision, not the panel, is the teaching: a presentation earns a lumbar puncture, the safety gate (papilloedema, any impaired consciousness, any focal sign) decides whether imaging comes first, and only a mass lesion or obstructive hydrocephalus, not raised pressure itself, forbids the needle. The study then returns pressure, colour, cells, protein, glucose read against a simultaneous blood glucose, and oligoclonal bands. Disease-specific profiles are deliberately absent. Three quantities that a reader may expect are omitted because they were not found in the sources read: the normal opening-pressure range, the CSF to serum glucose ratio, and the CSF lactate cut-off; the only pressure figure shown is the typical measured pressure in established idiopathic intracranial hypertension, which is neither a normal value nor a diagnostic criterion. Attributions per box (Kaufman's Clinical Neurology for Psychiatrists, Hutchison's Clinical Methods).

ASK	CHECK	SEND	PAIR
WHAT DIAGNOSIS COULD CSF CHANGE?	IS PUNCTURE SAFE NOW?	ONLY THE RELEVANT PANEL	RESULT WITH CONTEXT

5.4 Read the routine profile as a pattern

Normal values are the reference frame, not the final diagnosis. Normal CSF is clear. Its white cell count is **0 to 4 WBC/mL**, mostly lymphocytes; protein is **30 to 45 mg/dL**; and glucose is **60 to 100 mg/dL**. Glucose is never interpreted as an isolated CSF number: a simultaneous blood glucose sample is required. The source material establishes the pairing principle but not a numerical CSF-to-blood ratio, so the safe examination answer is to state the comparison without inventing a cut-off.

DOMAIN	NORMAL REFERENCE	INTERPRETIVE QUESTION
Colour	Clear	Is blood or pigment present?
White cells	Within the stated range, mostly lymphocytes	Is there pleocytosis, and which cell type predominates?
Protein	Within the stated concentration range	Does it rise with inflammation, or rise disproportionately to cells?
Glucose	Within the stated concentration range	How does it compare with simultaneous blood glucose?
Red cells	Interpret appearance and processing	Does blood represent an earlier subarachnoid bleed or procedural trauma?
Pressure	Measure when the clinical question requires it	Is pressure itself part of the suspected disorder?

The useful habit is to move through the same axes every time: appearance, pressure when indicated, cells, protein, glucose paired with blood, then targeted microbiology or immunology. A single abnormality rarely carries the whole interpretation. A cell count without the differential is incomplete; low CSF glucose without simultaneous blood glucose is context-poor; raised protein without attention to the cells misses a classic dissociation.

Interpretation is strongest when it separates observation from inference. “Lymphocyte-predominant pleocytosis with raised protein” describes the sample; “therefore this is a particular infection” goes beyond it and must be tested against the syndrome, timing, microscopy, molecular result, and culture. Conversely, a result within the routine reference profile does not answer a question that requires a specialised assay. The panel must match the hypothesis. This distinction prevents overcalling nonspecific abnormalities as well as false reassurance after an incomplete study.

Trend and coherence also matter, even when the page does not supply disease-specific numerical tables. Ask whether cells, protein, and glucose move in a mutually plausible direction, whether procedural contamination explains the appearance, and whether the result fits the neurological examination and imaging. If one element is strikingly discordant, reconsider the working

diagnosis, the sample, and the requested test before creating an elaborate explanation. CSF is evidence within a case formulation, not an autonomous diagnosis.

GOOD TO REMEMBER

Arrange the blood glucose at the same time as the puncture. CSF glucose must be compared with a simultaneous blood glucose concentration; a later or unpaired value weakens interpretation.

5.5 Infection and inflammation: order the panel, then recognise the pattern

When infection is the question, the minimum reasoning set is broader than "cells, protein, sugar." It includes white cell count, whether polymorphonuclear neutrophils or lymphocytes predominate, paired CSF and blood glucose, microscopy for microorganisms, polymerase chain reaction, and appropriate culture. Protein contributes to interpretation, and CSF lactate may be helpful. Microscopy may include Gram stain, staining for tuberculosis, and India ink for cryptococci. A numerical lactate threshold is not supplied here and should not be improvised.

- State the suspected infection so the laboratory can select appropriate microscopy, molecular analysis, and culture.
- Interpret the white cell count together with the differential rather than labelling every cellular response "infection."
- Pair CSF glucose with simultaneous blood glucose and treat protein as part of the pattern.
- Integrate microbiological evidence with the clinical syndrome and treatment timing.

Most infectious or inflammatory central nervous system illnesses produce **pleocytosis**. Protein tends to rise in parallel, while glucose falls to an abnormally low level in the characteristic generic pattern. This directional pattern is more defensible than memorising ungrounded disease-specific tables. It guides the next question, but overlap means it cannot replace culture, microscopy, molecular testing, or clinical judgement.

The classic exception is **albuminocytologic dissociation**: markedly elevated CSF protein with little or no increase in white cells, seen in Guillain-Barré syndrome. The teaching value is conceptual. A high protein result does not automatically imply a cellular inflammatory meningitis; the relationship between protein and cells is itself diagnostic information.

Suspected low-grade tuberculous meningitis may require repeated sampling on consecutive days before treatment to improve the chance of identifying and culturing mycobacteria. The verified source permits a finite repeat strategy, but the exact count is deliberately not printed here. This is less immediately urgent than an acute central nervous system infection, yet it remains an essential investigation when the syndrome warrants it.

5.6 Blood-stained CSF: traumatic tap or subarachnoid bleeding?

Blood introduced by the needle can falsely suggest subarachnoid haemorrhage. The laboratory distinction depends on centrifuging the bloody sample and examining the supernatant. A clear supernatant supports blood introduced during the procedure. A yellow supernatant indicates

xanthochromia, produced by degradation of haem after blood has been present in the CSF space, and therefore supports an earlier true bleed.

Timing is part of the test. When cranial imaging does not show subarachnoid blood but clinical suspicion remains, CSF examination should be delayed until at least **12 hours after the suspected bleed**. This interval allows red cells to lyse and haemoglobin to be metabolised into pigments detectable by spectroscopy or visible as xanthochromia. Sampling too early can make procedural blood and subarachnoid blood impossible to distinguish reliably.

TRAP **"Bloody CSF equals subarachnoid haemorrhage" is wrong.** Ask what the centrifuged supernatant shows and whether sufficient time elapsed for xanthochromia to develop.

This example captures a wider examination principle: pre-analytical details can determine whether a result is meaningful. The same observed abnormality may arise from disease or from the procedure. Interpretation therefore begins before the laboratory report, with the indication, timing, and quality of the sample.

5.7 Pressure and chronic inflammatory studies

Idiopathic intracranial hypertension illustrates why "raised intracranial pressure contraindicates lumbar puncture" is an unsafe oversimplification. After good-quality imaging has excluded a mass lesion, puncture is essential to measure pressure and analyse CSF. Established disease may show markedly raised measured pressure, with low protein, normal glucose, and no cells. The diagnostic threshold belongs to the dedicated idiopathic intracranial hypertension discussion and is not interchangeable with a typical measured pressure. The key distinction here is structural: raised pressure without a mass lesion does not create the same puncture hazard as a compartmental pressure gradient.

For suspected chronic inflammatory demyelination, CSF protein electrophoresis detecting **oligoclonal IgG** is the most sensitive evidence of chronic inflammatory central nervous system demyelination. In suspected multiple sclerosis, CSF is sent for oligoclonal bands and myelin basic protein. A very high cell count or very high protein may be incompatible with the working diagnosis and should trigger revision rather than forced confirmation.

CSF lymphocytosis can also accompany central nervous system vasculitis, neurosarcoidosis, Behçet disease, and sometimes systemic lupus erythematosus. These associations tell the psychiatrist when inflammatory disease should remain in the differential; they do not constitute a stand-alone diagnostic algorithm. The autoimmune and anti-NMDA receptor encephalitis antibody workup, along with the detailed neuroinfection panels, belongs in the Neuropsychiatry of systemic and neurological disease notes.

5.8 The neuropsychiatric workup: selective, not routine

A **dementia workup** does not routinely include lumbar puncture. It becomes appropriate when the presentation suggests an infectious cause, including neurosyphilis, subacute sclerosing panencephalitis, or Creutzfeldt-Jakob disease, or when suspected normal-pressure hydrocephalus makes pressure measurement and CSF withdrawal part of the diagnostic assessment. This is a

high-yield boundary: progressive cognitive symptoms alone do not justify a generic CSF panel, but a specific infectious, disease-directed, or pressure-related hypothesis may.

CLINICAL QUESTION	WHAT CSF CONTRIBUTES	BOUNDARY OF THIS SECTION
Infectious cognitive disorder	Disease-directed proteins, antibodies, retroviral markers, antigens, or polymerase-chain-reaction-detectable DNA	Order to the suspected organism or syndrome
Normal-pressure hydrocephalus	Pressure assessment and withdrawal of CSF	Interpret within the full clinical and imaging assessment
Progressive multifocal leukoencephalopathy in AIDS with a white-matter lesion	**JC virus CSF testing**; a positive result can obviate brain biopsy	A clear example of CSF replacing a more invasive test
Alzheimer disease biomarker question	Increased tau with decreased beta-amyloid and amyloid-beta peptide	The Dementias and neurocognitive disorders notes own the panel and its cut-offs
Autoimmune encephalitis or endemic neuroinfection	Targeted antibody or infection studies	The Neuropsychiatry of systemic and neurological disease notes own the algorithms and titres

Examples of disease-directed testing include protein associated with Creutzfeldt-Jakob disease, antimeasles antibodies in subacute sclerosing panencephalitis, retroviral markers in AIDS, and specific antigens or polymerase-chain-reaction-detectable DNA in other infections. These are not a licence to send every marker in every patient. Each test belongs to a prior clinical hypothesis, and the result must be interpreted within that disease pathway.

The same ownership principle prevents unnecessary duplication. Alzheimer disease CSF biomarkers may be named here by their direction of change, but their numerical interpretation belongs to the Dementias and neurocognitive disorders notes. Autoimmune encephalitis antibodies and detailed neuroinfection findings are indications for targeted CSF work, while their algorithms and titres belong to the Neuropsychiatry of systemic and neurological disease notes.

Before accepting the report, return to the original question. Confirm that the relevant measurement or assay was actually performed, that timing and treatment do not distort its meaning, and that the pattern fits the examination. If not, revise the hypothesis rather than stretching the result to fit it.

Memorise the sequence: define the question, exclude a dangerous pressure gradient or local infection, never delay urgent antimicrobial treatment, obtain pressure when it is part of the question, and interpret cells, protein, glucose, microbiology, and specialised studies as one clinical pattern.

6 · Soft neurological signs and neurological correlates of schizophrenia and developmental disorders

Soft neurological signs are not minor neurology. They are bedside observations that make a psychiatrist ask a more disciplined question: is the finding a focal deficit that requires localisation, a treatment-emergent movement disorder, or a non-localising sign of altered neurodevelopmental or integrative function? **Neurological soft signs** are clinically observable abnormalities that cannot be assigned to impairment of a particular brain region. "Soft" therefore describes their limited localising value, not their clinical importance or the amount of disability a patient experiences.

For psychiatry, the value of this lens is longitudinal and comparative. The signs can be present before drug exposure, may accompany cognitive and symptom dimensions, and can be encountered across psychotic and developmental presentations. They do not diagnose schizophrenia, ADHD, autism, or any other disorder in isolation. Instead, they sharpen formulation, improve the baseline against which later motor adverse effects are judged, and identify occasions when the ordinary psychiatric examination must widen into a careful neurological assessment.

6.1 What makes a sign "soft"?

A **hard neurological sign** has sufficient pattern, asymmetry, or associated deficit to support anatomical localisation. In contrast, soft signs are nonfocal abnormalities in the integration and execution of function. They are commonly observed through tasks involving coordination, sensory integration, developmental reflexes, balance, motor planning, and motor control. This is a functional description rather than a focal localisation.

RULE **"Soft" means non-localising, not benign.** A subtle, bilateral and non-lateralising performance abnormality can contribute to psychiatric formulation; a lateralised or gross abnormality must be read first as a possible focal neurological sign.

The distinction is easiest to understand when the same manoeuvre can be read in different ways. **Dysdiadochokinesia** means impaired rapid alternating movements. In a clearly lateralised presentation with intention tremor and other limb-coordination problems, it has localising significance for cerebellar hemisphere injury. A mild, bilateral irregularity without a localising pattern is instead an observation that may sit within a soft-sign assessment. The manoeuvre has not changed; its clinical meaning has changed with the total neurological pattern.

This is also why soft signs must never become a shortcut around a proper examination. A hard-sign finding should not be collapsed into an NSS score or explained away as "part of schizophrenia." The physical neurological examination and applied localisation are addressed in the Functional Neuroanatomy and Neurophysiology notes and the Stroke, TBI, delirium and focal brain syndromes notes.

PITFALL

Calling an abnormal movement “soft” before checking for asymmetry, weakness, sensory loss, altered reflexes, medication exposure, or a changing level of function is unsafe. Soft-sign language describes a pattern of non-localisation; it is not a permission to stop looking for neurological disease.

6.2 The domains that organise observation

The most useful teaching structure is the DSM-5-TR grouping: impairments in motor coordination, sensory integration, and sequencing of complex movements, with left-right confusion and disinhibition of associated movements included in the clinical picture. This grouping gives the examiner a way to notice patterns without pretending that any isolated task is diagnostic. Motor coordination asks whether an action is steady, timed, and appropriately graded. **Sensory integration** asks whether sensory information is being synthesised into an accurate percept or body schema. Sequencing asks whether the patient can organise a complex motor act in the required order. Associated-movement disinhibition concerns unwanted accompanying activity rather than the intended action alone.

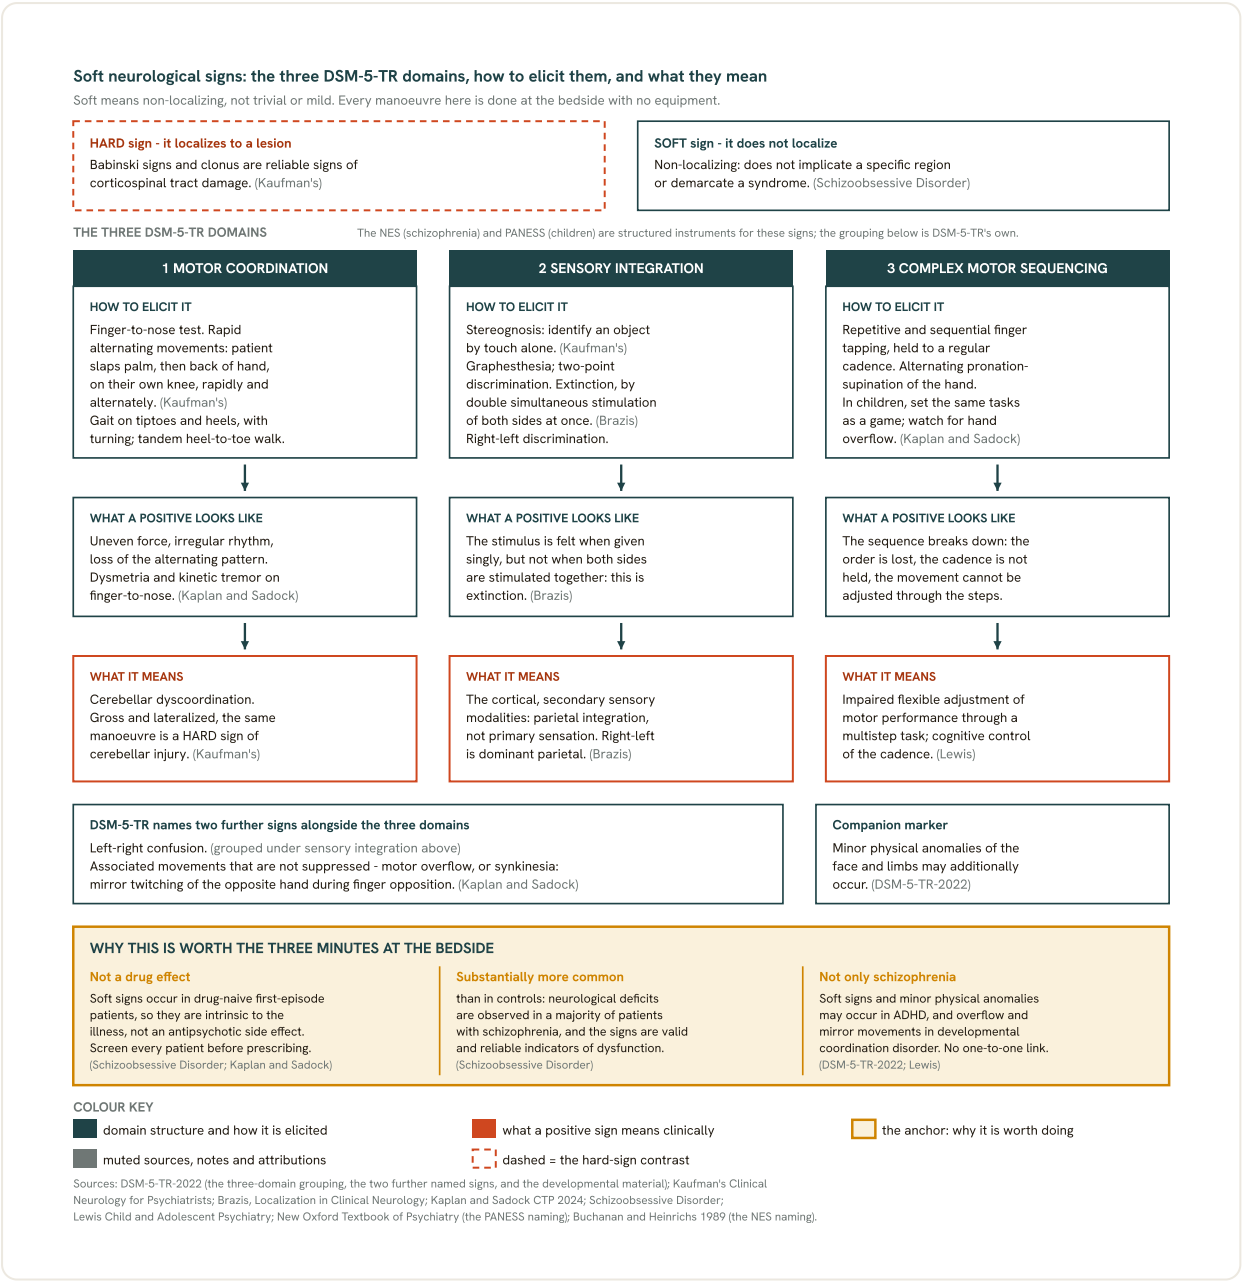


Fig 31.4 Soft neurological signs: the three DSM-5-TR domains, how each is elicited, and what a positive sign means. DSM-5-TR-2022 states that the soft signs common in schizophrenia comprise impairments in motor coordination, sensory integration and motor sequencing of complex movements, together with left-right confusion and associated movements that are not suppressed; minor physical anomalies of the face and limbs may additionally occur. The signs are called soft because they do not localize to a lesion, not because they are trivial or mild: the same manoeuvre that reads as a soft sign when it is subtle and bilateral reads as a hard sign when it is gross and lateralized. They are substantially more common in schizophrenia and in the developmental disorders than in controls, they are present in drug-naive first-episode patients, and every manoeuvre here is done at the bedside with no equipment. Grouping cited to DSM-5-TR-2022; elicitation and interpretation attributed per box (Kaufman's Clinical Neurology for Psychiatrists; Brazis, Localization in Clinical Neurology; Kaplan and Sadock CTP 2024; Schizoobsessive Disorder; Lewis Child and Adolescent Psychiatry).

CLINICAL DOMAIN	WHAT THE EXAMINER IS WATCHING FOR	INTERPRETIVE DISCIPLINE
Motor coordination	Balance, posture, precision and the smooth control of movement	Look for consistency, laterality and associated cerebellar or pyramidal signs before calling it soft.
Sensory integration	Integration of touch, position, body-side information and complex sensory perception	Ensure that primary sensation, comprehension and attention are adequate for the task.
Complex motor sequencing	Ordering, switching and sustaining a multistep motor programme	Separate an execution error from misunderstanding, fear, pain, fatigue, or severe psychopathology.
Associated-movement disinhibition	Unintended accompanying movement or overflow during a voluntary task	Describe what is seen and document baseline status before attributing later change to medication.

The accompanying figure is best used as a map for observation, not as a diagnostic algorithm. Within the motor domain, a rapid alternating-movement task can reveal irregular rhythm, uneven force, and failure to maintain alternation. Within the sensory domain, **stereognosis** is recognition of an object's form by touch. It is a cortical sensory modality, as are graphesthesia and related integrative tasks, and should not be mistaken for a basic test of light touch. Right-left confusion similarly belongs to an integrative task domain and should be interpreted alongside the rest of the neurological examination.

Extinction offers a useful contrast. On double simultaneous stimulation, the affected side is perceived when stimulated alone but not when both sides are stimulated together. This is a focal attentional-sensory phenomenon, not a generic soft sign to be casually scored. The lesson is practical: an NSS battery can direct attention to integrative function, but it cannot replace the clinician's duty to decide whether a result localises.

MNEMONIC

CIS

Use **C**oordination, **I**ntegration and **S**equencing as the bedside spine. Then look separately for associated-movement disinhibition. The mnemonic organises observation; it does not convert the domains into diagnostic criteria.

6.3 Primitive reflexes and motor overflow

Primitive reflexes require particular care because their presence can be discussed under soft-sign traditions, yet the finding may carry a harder neurological implication in the appropriate context. **Frontal release reflexes** include snout, suck, rooting, jaw jerk, palmomental and grasp responses, most of which are normally present in infancy. Their appearance in an adult is abnormal. Several such reflexes, interpreted with the history and the remainder of the examination, can suggest

congenital cerebral injury, frontal lobe damage, or neurodegenerative illness. They should never be treated as a diagnostic signature of a functional psychiatric syndrome.

For the psychiatrist, their importance is that pathological reflexes can support an organic basis for personality change and correlate to some degree with cognitive impairment. The relevant frontal-system anatomy belongs in the Functional Neuroanatomy and Neurophysiology notes; here, the examination lesson is to record the sign, its distribution, and its neurological context rather than using it as decorative neuropsychiatric jargon.

Synkinesia, or motor overflow, denotes unintentional movement accompanying a voluntary action. Mirror movements, contralateral finger twitching during a finger-opposition task, and facial grimacing during hand movements illustrate the phenomenon. In children, overflow can be read as a developmental immaturity; in adults, the same observation needs to be set against medication history, neurological examination, and course. “Disinhibition of associated movements” is a descriptive observation, not a licence to infer a particular lesion.

6.4 Schizophrenia: correlates, course and meaning

NSS are common in schizophrenia, but they are not specific to it. Their significance lies in convergence with a developmental reading of illness, not in diagnostic exclusivity. They have been observed in **dopamine-antagonist medication-naive first-episode psychosis** and in most people with chronic schizophrenia. This makes a purely iatrogenic explanation inadequate, while still leaving room for medication to add new extrapyramidal phenomena later in the course.

Evidence also links NSS with cognitive deficits and with negative and disorganised symptoms. A further association has been reported with cognitive and motivation-expression deficits and with positive psychopathology resistant to initial antipsychotic treatment. These are correlates, not a severity scale that determines treatment choice. The clinician should therefore avoid using a soft-sign burden to make a prognostic declaration to a patient or family. It is better used to enrich the baseline description, direct collateral developmental history, and make monitoring for treatment-emergent movement change more defensible.

There is a trait argument and a state argument, and neither should be caricatured. Soft signs have been identified in people at high risk, in monozygotic co-twins discordant for schizophrenia, and in first-degree relatives, supporting their potential as **trait markers**. Prospective developmental evidence is also striking: children who later developed schizophrenia had **more than one soft neurological sign** and poorer motor-skill performance at several childhood assessments. Yet some motor and cortical signs may improve as psychopathology improves, whereas harder signs can appear more static and trait-like. The careful conclusion is that the construct contains both state-sensitive and trait-like elements.

This pattern fits the broader association of neurological signs with cognitive dysfunction, developmental anomaly, and diffuse brain pathology. It should be named, then bounded: the full neurodevelopmental account of schizophrenia belongs in the Schizophrenia: aetiology, phenomenology, classification and course notes. Reported neural-circuit and structural correlations are mechanistic correlates, not a bedside imaging indication.

CLINICAL QUESTION	WHAT NSS CAN CONTRIBUTE	WHAT NSS CANNOT ESTABLISH
Was movement present before treatment?	A documented baseline for later comparison	That a later movement is necessarily unrelated to medication
Is there a neurodevelopmental signal?	A prompt to obtain developmental and family history	A diagnosis of schizophrenia or aetiological proof
Is cognition relevant to current disability?	A reason to assess cognition carefully and functionally	A substitute for formal cognitive assessment
Could this be focal disease?	A cue to inspect for localising features	A reason to ignore a focal or progressive deficit

The cognitive screen and its formal interpretation are covered in the Psychometry, Assessment and Descriptive Psychopathology notes. The acute adverse-effect framework belongs in the Schizophrenia: management, antipsychotics and adverse effects notes. The present task is narrower and clinically potent: identify what was present before prescribing, describe it precisely, and refuse false certainty about causation.

NONLOCALISING	CIS	BASELINE	INTRINSIC
DEFINING FEATURE	CLINICAL SPINE	BEFORE PRESCRIBING	NOT SIMPLY DRUG EFFECT

6.5 Structured assessment and the baseline motor record

The **Neurological Evaluation Scale** is a structured instrument for assessing neurological signs in schizophrenia. It contains **26 items**, organised into motor coordination, sensory integration, sequencing of complex motor acts, and a residual "others" grouping. Each item is rated on an ordinal **0 to 2** scale using standardised instructions, with higher scores indicating greater disruption. Its psychometric literature should be read with attention to the instrument's stated scope and the licence boundary.

The NES is useful for standardising observation and for research-quality description. It is not a neurological consultation in a booklet, and this chapter does not reproduce its protected items, wording, or scoring anchors. In ordinary practice, a structured screen should be embedded within the general motor examination: gait, station and balance, abnormal movement at rest and with arms outstretched, coordination, tone, and strength. Finger-to-nose and rapid alternating movements are standard cerebellar manoeuvres; finger-to-nose also samples proximal arm movement, kinetic tremor, and dysmetria.

GOOD TO REMEMBER

Record spontaneous movements and soft signs before the first psychiatric drug exposure whenever feasible. Abnormal movements can precede treatment and have a broad differential, including congenital, developmental, neurodegenerative and acquired cerebral conditions, as well as normal ageing. A later adverse-effect judgement is only as good as the baseline record.

Baseline signs may also be relevant when monitoring for treatment-emergent extrapyramidal adverse effects. This does not make NSS a deterministic prediction tool, but it is a strong practical

reason to document the motor examination rather than retrospectively reconstruct it after an adverse effect occurs. Careful appraisal of movement belongs to general psychiatric practice because emotional, cognitive and motor disorders involve interrelated basal-ganglia and frontal-subcortical circuits.

6.6 Developmental interpretation in children

In developmental psychiatry, a soft sign is often best understood as delay in a normally acquired achievement: integration of motility, smooth and accurate motor sequencing, or auditory-visual integration. This language avoids prematurely labelling a child with a neurological abnormality when the observed phenomenon may represent immaturity. Its value depends on developmental context, functional impact, and whether the observation persists beyond what is expected for the child.

Associations with psychopathology have been demonstrated when **at least 3 soft neurological signs** are present in children older than **7 years**. The age qualifier matters: earlier in development, many signs may represent delayed achievements rather than persisting dysfunction. This is a threshold for an association, not a diagnostic criterion and not an invitation to count minor behaviours without a developmental examination.

Assessment can be made collaborative and game-like. Useful observations include ball throwing, walking on a line, heel-to-toe walking, overflow of the hands while walking on the outside of the feet, repetitive or sequential finger tapping, and alternating pronation-supination. The child should understand the task and have an opportunity to practise. A poor performance can arise from anxiety, inattention, oppositionality, pain, unfamiliarity, language difficulty or reduced comprehension, so the examiner must distinguish a stable motor pattern from situational non-cooperation.

- Observe quality, not merely completion: rhythm, bilateral symmetry, sequencing and overflow often provide more information than success or failure.
- Compare across tasks: a single awkward action is weak evidence, whereas a convergent pattern across coordination, integration and sequencing is more meaningful.
- Link the finding to daily function: handwriting, play, self-care, classroom demands and social confidence determine its practical significance.

In developmental coordination disorder, mirror movements and choreiform movements of unsupported limbs are described as neurodevelopmental immaturities or soft signs rather than neurological abnormalities, and their diagnostic role remains unclear. Developmental coordination disorder occurs in **5 to 8 per cent** of children aged **5 to 11 years**. The crucial clinical distinction is that overflow may support a developmental formulation but does not independently make the diagnosis.

6.7 ADHD, communication disorders and broader developmental correlates

ADHD has no diagnostic physical sign, though subtle motor delays, NSS and minor physical anomalies may occur. **Minor physical anomalies** described in this context include

hypertelorism, a highly arched palate and low-set ears. Marked co-occurring clumsiness or motor delay should prompt separate consideration of developmental coordination disorder rather than being absorbed into ADHD. Soft signs or minor anomalies may modestly increase confidence in an ADHD diagnosis, but there is no one-to-one relation.

The proposed reading of ADHD soft signs concerns cognitive control of motor inhibition and excitation during repetitive actions, flexible adjustment in multistep tasks, and integration of proprioceptive input and body-position sense in gait and balance. A reported association with reduced bilateral frontal cerebral blood-flow measures does not make functional imaging part of routine ADHD workup. The finding is a correlate, not a test to order.

Similar observations have been described in speech sound disorder and language disorder, alongside subtle motor-skill deficits and mixed cerebral dominance. These observations argue for a broad developmental review where function is impaired, not for the false claim that an NSS profile can identify a single psychiatric diagnosis.

TRAP **Do not turn a developmental correlate into a diagnostic sign.** Clumsiness, overflow, minor anomalies and soft signs can strengthen the decision to examine development more fully, but they do not replace syndrome-level history, impairment assessment or differential diagnosis.

6.8 A practical psychiatric approach

LOCUS is the routine outpatient, emergency or inpatient assessment where movement and neurological history should be captured before treatment alters the picture. FOCUS is the current motor phenotype, its onset and course, associated cognition or behaviour, and any localising or progressive features. MODUS is careful observation, a focused neurological examination, structured documentation when useful, collateral developmental history, and escalation for neurological evaluation when the pattern is focal, progressive, or otherwise concerning.

For a person presenting with schizophrenia, an examination focused on the extrapyramidal system should be part of routine assessment. This is not an optional neuropsychiatry flourish. It protects against misattributing a pre-existing movement abnormality to treatment, provides a defensible baseline, and may identify an alternative neurological explanation for apparent psychiatric change. In a child, the same principle means observing movement in an age-appropriate context and making function, not a checklist total, the centre of interpretation.

A soft neurological sign is a non-localising clue to integrative or developmental dysfunction; document it, contextualise it, and never let it obscure a focal neurological deficit or a medication-related movement disorder.

7 · Disorders of consciousness, coma, the vegetative and minimally conscious states and brain death

The marks lie at the boundaries. Separate arousal from awareness, coma from wakeful unresponsiveness, purposeful behaviour from reflex movement, and severe impairment from

death. A wrong label may conceal reversible neurological disease, deny a conscious person communication, or confuse a family during critical care.

Use an anatomical and behavioural approach: test arousal, seek reproducible evidence of awareness, examine the brainstem rostrocaudally, and describe observations before naming a syndrome. Higher-function testing belongs to the arousable patient. The descriptive ladder is covered in the Psychometry, Assessment and Descriptive Psychopathology notes, while delirium belongs to the Stroke, TBI, delirium and focal brain syndromes notes.

7.1 The architecture of consciousness

Consciousness depends on interaction between the ascending reticular activating system and the cerebral hemispheres. This gives the clinician separable axes: arousal, supported by the activating system, and awareness, expressed through organized hemispheric function. Alertness is necessary for ordinary awareness but does not guarantee cognition, emotional capacity, or meaningful engagement. Conversely, a patient may retain an intact mind while motor output is almost completely abolished.

The ascending reticular activating system occupies the paramedian tegmentum of the posterior pons and midbrain and extends toward the posterior hypothalamus and thalamic reticular formation. Diffuse thalamocortical projections regulate cortical activity. Ocular motor structures lie within this system, explaining why eye and pupil findings localise brainstem dysfunction. Sedatives may depress the same synaptic network that structural disease can damage.

RULE **Coma requires failure of the activating system or failure of both cerebral hemispheres.** A unilateral hemispheric lesion is therefore an exceptional and usually transient explanation unless it produces secondary brainstem compromise.

Clinical neurology describes a descending sequence of **alert, lethargic, stuporous, and comatose**. A lethargic patient appears asleep but can be brought temporarily to an awake state by sufficient stimulation, remaining inattentive, disoriented, and cognitively impaired when roused. A stuporous patient remains unarousable and gives only rudimentary motor or verbal responses. In coma, the eyes remain closed, limb movement is reflexive, verbal response is absent or minimal, and there is no meaningful environmental interaction. Coma is usually transitional, lasting from hours to weeks before recovery, death, or evolution into another disorder of consciousness.

7.2 Bedside examination of coma

Build a pattern from pupils, spontaneous and reflex eye movements, motor symmetry, posturing, and respiration while assessing reversible toxic and metabolic causes. The light reflex resists metabolic dysfunction, so unilateral abnormality strongly suggests a structural midbrain or oculomotor lesion. Drugs, hypothermia, and acute anoxia can invalidate this shortcut.

- **Pupils:** metabolic or bilateral diencephalic dysfunction usually leaves small, reactive pupils. Tectal or pretectal midbrain dysfunction produces midsized or slightly enlarged pupils without a light response but with a ciliospinal response. Tegmental midbrain dysfunction

produces midsized, often unequal pupils lacking both responses. Pontine dysfunction causes small pupils, while pontine haemorrhage may produce pinpoint pupils that still constrict under magnification.

- **Herniation:** lateral transtentorial compression of the oculomotor nerve produces an ipsilateral, widely dilated pupil with a sluggish or absent light response, the Hutchinson pupil. If hemiparesis appears on the same side as the dilated pupil, the Kernohan notch phenomenon is a possible false-localising explanation.
- **Eye movements:** as coma deepens, roving movements disappear before the oculocephalic response, and the caloric response is lost last. Pupillary reactivity may persist after eye movements are no longer elicitable in metabolic coma.
- **Respiration:** Cheyne-Stokes respiration is associated particularly with bilateral thalamic dysfunction but may occur with descending-pathway lesions extending to the upper pons. Apneustic breathing points toward the lower pontine lateral tegmentum, cluster breathing toward low pontine or high medullary dysfunction, and ataxic breathing toward dorsomedial medullary damage and impending respiratory failure. Metabolic disturbance can imitate these patterns.

Oculovestibular testing supplies a stronger stimulus than the oculocephalic manoeuvre. After confirming an intact tympanic membrane and elevating the head by 30 degrees, 50 mL of ice water is instilled over 30 seconds. This is a specialist bedside procedure, not a casual test, because technique and ear pathology determine whether an absent response is interpretable.

MNEMONIC

Pupils, Ocular, Corneal, Gag, Motor

The brainstem-reflex sweep in coma follows **P**upillary light response, **O**cular responses, **C**orneal reflexes, **G**ag and cough, and **M**otor response. Read the findings as a linked rostrocaudal pattern, never as isolated tricks.

7.3 Motor response and posturing

Decorticate posturing consists of shoulder and arm adduction, elbow flexion, and wrist pronation and flexion, with extension of the leg. With a unilateral hemispheric lesion it is contralateral. **Decerebrate posturing** consists of upper-limb extension and pronation with forceful plantar flexion, and may follow severe metabolic injury or upper brainstem dysfunction. More caudal progression may produce arm extension with weak leg flexion in pontine tegmental damage and total flaccidity with medullary involvement.

TRAP **Posturing does not prove a structural lesion.** Toxic-metabolic coma commonly produces both flexor and extensor patterns. Structural patterns are often asymmetric, but the posture name alone must not be converted into an anatomical diagnosis.

Motor symmetry has greater discriminating value. Marked asymmetry argues against purely metabolic encephalopathy. Tremor, asterixis, myoclonus, and seizures favour toxic-metabolic

dysfunction, although metabolic and structural disease may coexist. Asterixis may disappear as coma deepens. Document the stimulus, best response on each side, spontaneous movement, and serial change.

7.4 Structured assessment of responsiveness

The **Glasgow Coma Scale** records eye opening, verbal response, and best motor response. It is a compact description of observable neurological function, not a diagnosis and not a substitute for the brainstem examination. Report the components as well as the total because identical totals can represent clinically different response patterns.

COMPONENT	RESPONSE	SCORE
Eye opening	Never	1
	To pain	2
	To verbal stimuli	3
	Spontaneously	4
Best verbal response	None	1
	Incomprehensible sounds	2
	Inappropriate words	3
	Disoriented and converses	4
	Oriented and converses	5
Best motor response	None	1
	Extension, decerebrate rigidity	2
	Flexion, decorticate rigidity	3
	Flexion withdrawal	4
	Patient localizes pain	5
	Patient obeys	6

The total ranges from **3 to 15**, with lower totals indicating poorer neurological function. In the traumatic brain injury severity bands used by Kaufman, 3-8 denotes severe injury or coma, 9-12 moderate injury, and 13-15 mild injury. These are trauma-severity bands, not a universal taxonomy of consciousness, and they must not be merged with alternative bands that classify a different construct.

E	V	M	E + V + M
EYE OPENING	VERBAL RESPONSE	MOTOR RESPONSE	REPORT COMPONENTS

The scale correlates closely with outcome after major head trauma but poorly after minor trauma. On the first day after traumatic brain injury, **90%** of patients at the minimum score have a fatal outcome, and most survivors do not regain consciousness. Yet that score cannot determine death.

Cerebral hypoxia and intubation limit interpretation, and the scale does not belong in a mental status examination for dementia or circumscribed cognitive deficit: it grades responsiveness, not cognition.

PITFALL

Writing only a total hides clinically important information and may create false precision when the verbal component is untestable. Record the observed component pattern and the reason for any untestable response.

A low Glasgow Coma Scale score describes impaired responsiveness; it never, by itself, declares death.

Teasdale G, Jennett B. Assessment of coma and impaired consciousness: a practical scale. Lancet 1974;2(7872):81-84. The scale reproduced here is acknowledged as copyright of Glasgow University and is cleared only for the free, non-commercial resident PDF. This clearance does not extend to related prognostic charts, imputation tools, or commercial channels.

The accompanying figure consolidates the bedside distinctions described here.

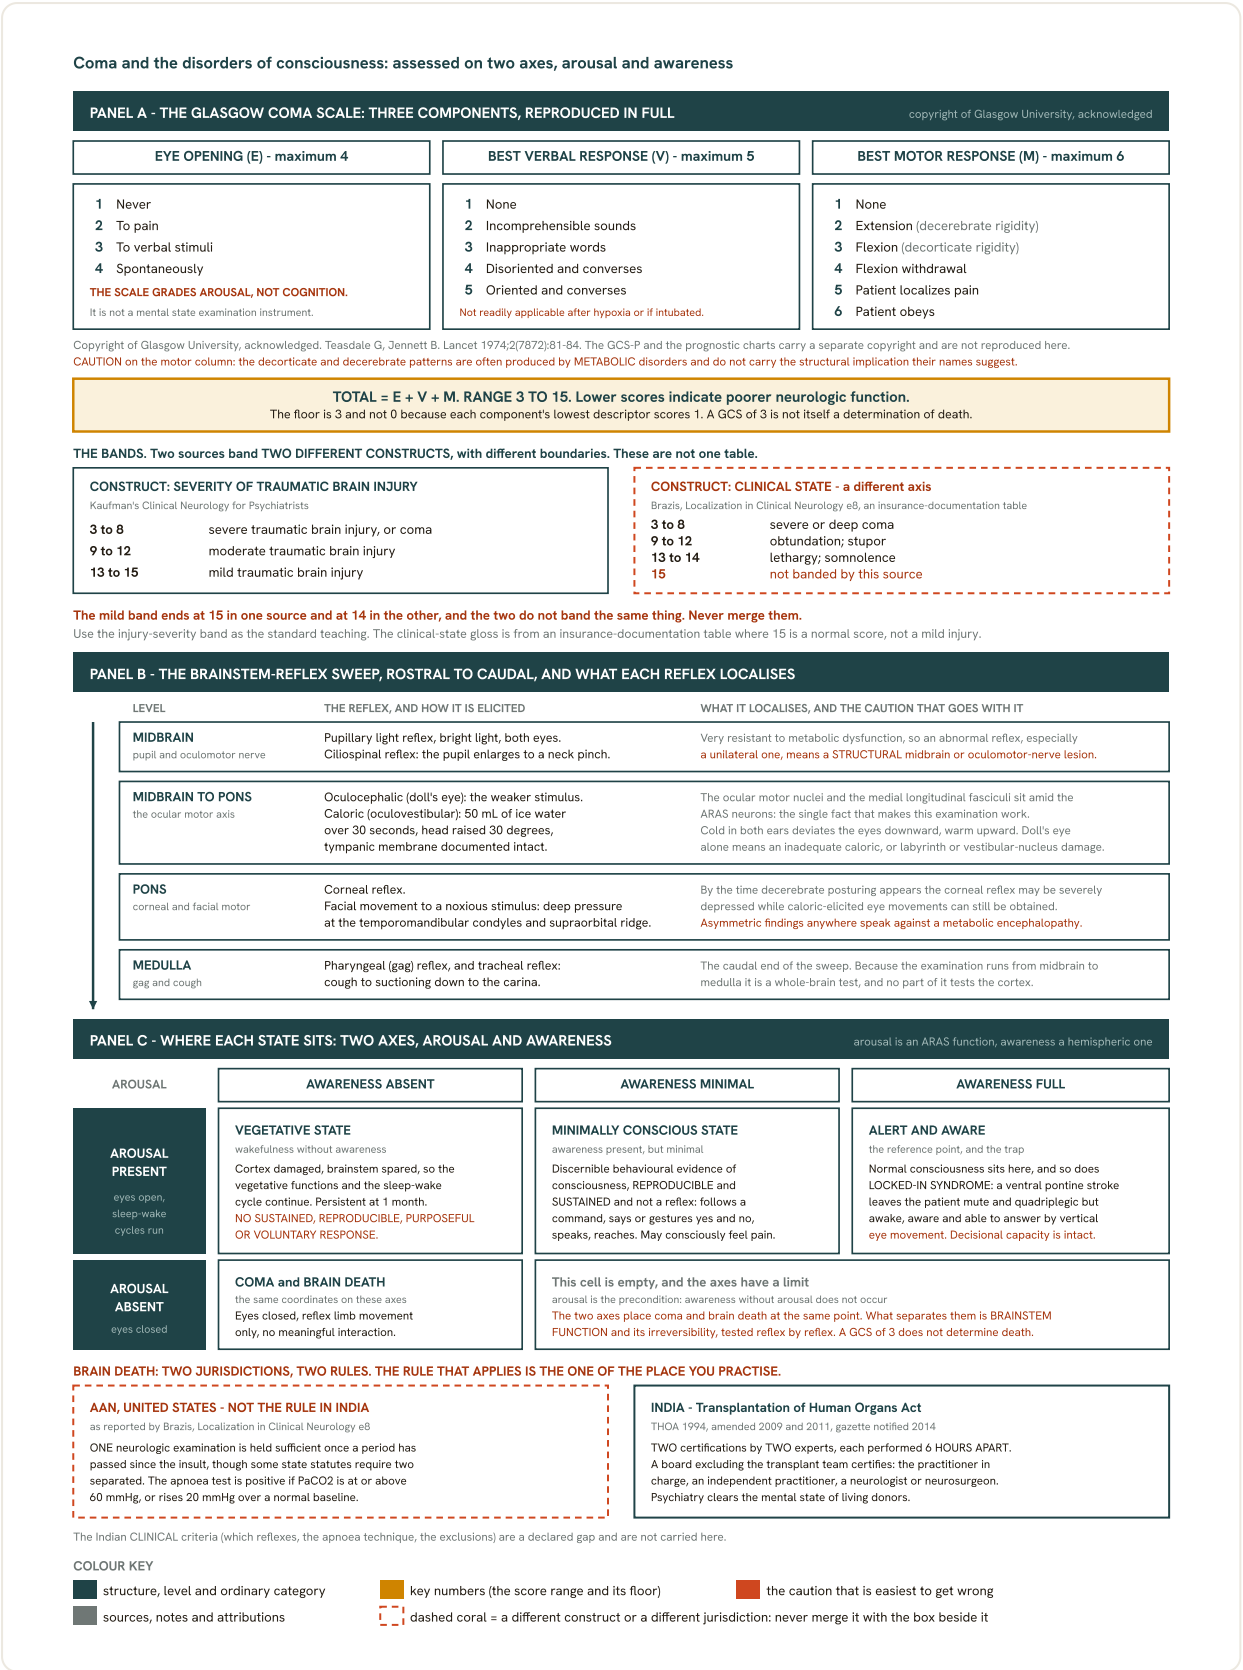


Fig 31.5 Coma and the disorders of consciousness: the Glasgow Coma Scale, the brainstem-reflex sweep, and where each state sits. Consciousness is assessed on two axes, arousal and awareness: coma means either the ARAS is down or both hemispheres are. The Glasgow Coma Scale is the copyright of Glasgow University, acknowledged on the plate, where the licence conditions sit. Brazis prefers 'unresponsive wakefulness syndrome', on the ground that responses are all the examiner can evaluate. (Kaufman's Clinical Neurology for Psychiatrists; Brazis, Localization in Clinical Neurology e8; RC Jiloha, Forensic Psychiatry: An Indian Perspective.)

7.5 Vegetative and minimally conscious states

Vegetative state is wakefulness without awareness: widespread cortical or immediately subcortical damage coexists with relative brainstem preservation. The eyes open, sleep-wake cycling continues, and vital vegetative functions persist, but there is no evidence of awareness of self or environment, language comprehension or expression, or sustained, reproducible, purposeful or voluntary behaviour. Cranial and spinal reflexes vary, and incontinence is typical.

The word persistent is temporal rather than synonymous with permanent. It applies when the vegetative state remains for **at least 1 month** after acute traumatic or nontraumatic injury, or for the same duration in degenerative, metabolic, or developmental disease. Prognosis depends strongly on cause and duration. Kaufman gives categorical irrecoverability thresholds of 1 year after traumatic injury and 3 months after neurodegenerative illness, whereas Brazis frames recovery after comparable durations as unlikely or rare rather than impossible. The conflict should be stated, not flattened into certainty.

Minimally conscious state retains reproducible, sustained, non-reflex behavioural evidence of consciousness. This may be command-following, a verbal or gestural yes-no response even if inaccurate, intelligible speech, appropriate smiling or crying, an answer by speech or gesture, or purposeful reaching and manipulation. Such patients may consciously perceive pain and generally have a more favourable outcome than patients in a vegetative state. Organized cortical activity may be evident on EEG or functional imaging in the minimally conscious state when meaningful activity is absent in the vegetative state.

STATE	AROUSAL	AWARENESS AND BEHAVIOUR	KEY SEPARATION
Coma	Absent	No meaningful interaction	Eyes remain closed
Vegetative state	Present	No sustained, reproducible purposeful response	Wakefulness without awareness
Minimally conscious state	Present but impaired	Reproducible, sustained, non-reflex evidence	Behaviour proves at least intermittent awareness
Locked-in syndrome	Present	Awareness and decisional capacity intact	Output failure, not impaired consciousness
Death by neurological criteria	Absent	Irreversible cessation of all functions of the entire brain	A formal jurisdiction-bound determination

The diagnosis remains epistemically difficult because motor output is the examiner's access to awareness, yet motor response is precisely what severe injury may abolish. Some patients meeting behavioural criteria for chronic vegetative state have shown command-following or awareness on functional neuroimaging. For this reason, unresponsive wakefulness syndrome or persistent unresponsiveness may be preferable language. Imaging may challenge the bedside inference, but structural damage alone does not settle awareness.

7.6 Locked-in syndrome and akinetic mutism

Locked-in syndrome is the inverse of vegetative state. The patient is mute and quadriplegic, yet awake, aware, cognitively intact, and capable of decisions. Eye or eyelid movement permits communication. A typical ventral pontomedullary lesion interrupts bulbar and corticospinal output while sparing the dorsal activating system, hemispheres, and ocular pathways. The EEG is relatively normal.

> **GOOD TO REMEMBER**
>
> In an apparently unresponsive patient with open eyes, ask for vertical eye movement or eyelid movement to command. Correct responses permit a communication code and a formal assessment of cognition and capacity. Horizontal gaze may be impaired by the pontine lesion.

Severe peripheral weakness can create the same de-efferented appearance. Myasthenia gravis, amyotrophic lateral sclerosis, and Guillain-Barre syndrome may abolish cranial and limb movement; a complete locked-in state can even lose pupillary responses and be mistaken for anoxic brain damage. A normal or mildly slow EEG is a crucial clue that cerebral function persists. The ethical consequence is immediate: a conscious patient may express treatment preferences once a reliable communication method is established.

Akinetic mutism describes a seemingly awake patient who is silent and motionless, with eyes that may track movement but little evidence that attention or interest can be engaged. Descending motor tract signs are limited, while frontal release signs may occur. Bilateral anterior cingulate or frontal lesions, diencephalomesencephalic reticular dysfunction, and lesions of the globus pallidus or hypothalamus are recognised substrates. Autonomic arousal with tachycardia and sweating during discomfort may resemble catatonia and does not discriminate. Frontal release or corticospinal signs and slow-wave EEG abnormalities favour akinetic mutism; a normal or desynchronised low-voltage fast EEG favours catatonia.

7.7 Functional, ictal and autoimmune mimics

Do not diagnose function by theatricality. In functional or dissociative unresponsiveness, the eyelids may be held forcibly shut against opening, or the patient may stare with quick blinks. Pupils are normally positioned and reactive, tone and reflexes are normal, and the oculocephalic response may be random or absent. The decisive discriminator is caloric stimulation: classic vestibular nystagmus with a preserved quick phase requires functioning frontal eye fields and therefore demonstrates cortical activity. True coma lacks that quick phase.

Roving eye movements of light coma cannot be produced voluntarily and therefore argue against feigning. Forced downward gaze has been reported in feigned coma but is not specific because sedative-induced coma may show it during caloric testing. Recurrent functional unresponsiveness is often misdiagnosed as epilepsy or migraine, replacing formulation with repeated disease-labelled treatment.

Nonconvulsive status epilepticus may resemble metabolic encephalopathy without clear focal signs. Passively elevate the lids and look for small-amplitude, mainly vertical eye movements,

then inspect for subtle intermittent focal or multifocal rhythmic movement. Suspicion warrants prompt EEG; the broader seizure assessment belongs to the Epilepsy and psychiatry notes.

Anti-NMDA receptor encephalitis is a critical psychiatric trap. A psychiatric presentation may be followed by choreic or other movement disorder, respiratory insufficiency, and coma. Akinetic-mutism-like or catatonia-like states may alternate with agitation, while resistance to eye opening and absent response to pain can misleadingly appear dissociative. The combination of psychiatric symptoms and apparently functional examination findings must therefore increase, not reduce, concern for this treatable autoimmune disorder. Its full investigation belongs to the Neuropsychiatry of systemic and neurological disease notes.

7.8 Death by neurological criteria

Death by neurological criteria is clearer clinical language than brain death because it avoids implying an intermediate state followed by a later, more real death when ventilation is withdrawn. The diagnosis requires irreversible cessation of all functions of the entire brain. It is categorically different from coma, vegetative state, minimally conscious state, and locked-in syndrome.

The following clinical framework is the American Academy of Neurology position for the United States, not the Indian clinical standard. Before examination, clinicians establish an irreversible proximate cause, exclude central nervous system depressant drugs and neuromuscular blockade, exclude severe electrolyte, acid-base, or endocrine disturbance, ensure core temperature is above 36 degrees Celsius, and ensure systolic blood pressure is at least 100 mmHg.

The neurological assessment requires coma, absence of brainstem reflexes, and apnoea. The reflex examination demonstrates absent bilateral pupillary light responses, absent oculocephalic and oculovestibular responses, absent corneal responses, absent facial movement to noxious stimulation, and absent gag and cough. Spinally mediated movements do not refute the diagnosis when every brainstem-mediated response is absent.

Under the same United States framework, the apnoea test evaluates respiratory drive using a carbon dioxide challenge after physiological prerequisites are met and oxygenation is maintained. It supports the diagnosis when respiratory movement remains absent and arterial carbon dioxide reaches at least **60 mmHg**, or rises by 20 mmHg above a normal baseline. This is an intensive-care procedure with explicit preparation and abort conditions, not an examination-room manoeuvre.

Ancillary testing is used when clinical assessment cannot be completed or interpreted. Hypothermia, drug overdose, and metabolic derangement can mimic electrocerebral silence; silence alone does not prove irreversibility, while persisting EEG activity refutes the diagnosis. Fixed pupils after acute neurosurgical injury also do not prove irreversibility: among 40 patients, 25% recovered function, although none recovered after more than 6 hours of fixed pupils.

Spinal automatisms may persist after total irreversible brain destruction. These include withdrawal, stretch reflexes, toe or abdominal movements, respiratory-like movements, and the

Lazarus sign. Such automatisms occur in up to 40% of heart-beating cadavers, usually within the first 24 hours. Their apparent purposefulness can distress families and junior staff; the decisive point is that no facial, gag, cough, or other brainstem-mediated response is present.

7.9 Indian certification and psychiatric liaison

IN INDIA

The Transplantation of Human Organs Act, 1994, amended in 2009 and 2011 and gazette notified in 2014, recognises brain death as a definite form of death. Certification requires two certifications by two experts, performed 6 hours apart, and certifiers must not belong to the transplant team. These are certification rules. The allowed evidence does not supply Indian clinical testing criteria, so the United States criteria above must not be substituted for local statutory, state, and institutional requirements.

LOCUS: care occurs in the emergency department and intensive-care unit, with later work at the bedside family meeting and within the transplant pathway. **FOCUS:** acutely, protect against a false diagnosis by establishing irreversibility and excluding confounders; subsequently, provide an unambiguous explanation of death, prepare relatives for spinal movements, and separate the death determination from the donation decision. **MODUS:** the work is clinical, procedural, statutory, communicative, ethical, and social. The governing frames are the American position only where its jurisdiction applies and the Indian statutory framework.

For the psychiatrist, capacity problems sit on opposite sides of this boundary. A locked-in patient needs a communication channel before capacity is judged. In transplantation practice, psychiatry clearance is required for the mental state of living donors. In brain death, the primary team raises organ donation and the next of kin may authorise it. Psychiatry identifies impaired understanding, coercion, conflict, or psychopathology and keeps death communication separate from donation counselling.

Families often interpret open eyes, reflex orienting, posturing, or spinal automatisms as awareness or suffering. Clear language should acknowledge what they have seen, state which responses are reflexive, explain which findings would demonstrate awareness, and avoid implying that withdrawal of ventilation causes death. That precision is both neurologically correct and psychologically protective.

Headache and the psychiatric interface

8 · Migraine: psychiatric comorbidity (depression, anxiety, bipolar disorder) and shared mechanisms

Migraine is a neurological disorder with a psychiatric interface, not a personality diagnosis. For the psychiatrist, the important task is to recognise a genuine migraine syndrome, identify comorbid mood and anxiety symptoms without collapsing them into the headache, and decide whether a temporal relationship is evidence of causation or simply coexistence. The distinction

matters because symptoms may be present during an attack, within an aura, between attacks, or as an independent psychiatric disorder. Each setting calls for a different formulation.

The International Headache Society categorisation places migraine among primary headaches. That framing is useful in psychiatric practice: the diagnosis rests on the characteristic clinical syndrome, whereas examination and laboratory findings are generally normal. It does not mean that the suffering is minor. Recurrent pain, sensory sensitivity, anticipatory fear, disrupted functioning and repeated health-care contact can all shape mood and behaviour. Conversely, depression, anxiety and bipolar-spectrum symptoms can alter help-seeking, adherence, sleep, substance use and the interpretation of bodily symptoms. The consultation must therefore hold both sides of the interface in view.

8.1 Establish the migraine phenotype before attributing psychiatric meaning

A psychiatric formulation starts with phenomenology. The **ICHD-3 migraine without aura criteria** require a recurrent attack pattern, a headache lasting **4-72 hours** when untreated or unsuccessfully treated, characteristic pain features, associated nausea or vomiting or light and sound sensitivity, and exclusion of a better explanation. This is not an invitation to diagnose every severe headache as migraine. It is a reminder that the syndrome has a recognisable structure, and that the history should test that structure before depression or anxiety is made the explanatory centre of the case.

The distinction is particularly important when patients describe neurologic symptoms as “panic,” “dissociation,” or “hallucinations.” In **migraine with aura**, the neurologic symptoms are fully reversible and may be visual, sensory or dysphasic, without motor weakness in the typical-aura description. A symptom may develop gradually and each aura symptom has a bounded duration of **5-60 minutes**. The headache may begin during the aura or follow it. This temporal sequence is often more diagnostically useful than the emotional language used to describe it.

Ask for a chronological account in ordinary language: what was noticed first, how it changed, whether a positive phenomenon appeared alongside a loss of function, what followed, and whether the same sequence has recurred. This protects the patient from an erroneous psychiatric label and protects the psychiatrist from missing a neurological syndrome. It also avoids a false binary in which a patient is assumed to have either migraine or emotional distress. An attack can be neurologically typical and psychologically frightening at the same time.

The history is also where the psychiatrist decides what needs parallel assessment. A stereotyped recurrent sequence makes a migraine formulation more coherent, but it does not abolish the need to examine a new or changed presentation in its clinical context. Conversely, a normal neurological examination between familiar attacks should not be used to imply that the pain, disability or comorbid emotional symptoms are unreal. The psychiatric task is interpretive: identify the syndrome, the timing of its effects and the patient’s response to them, rather than forcing all symptoms into a single category.

RULE **Diagnose the headache syndrome and the psychiatric syndrome on their own evidence before deciding how they are related.** Temporal proximity alone can support a hypothesis, but it cannot settle causation.

8.2 Depression: a reciprocal association, not a simple reaction to pain

Migraine-depression comorbidity is clinically substantial and appears reciprocal. Major depression increases subsequent migraine risk, while migraine is associated with an increase in major-depression risk of **up to 4-fold**. The practical implication is that neither diagnosis should automatically be treated as secondary to the other. Pain can worsen mood, but the observed association is not exhausted by understandable demoralisation in a person with recurrent pain.

~40%	~60%	up to 4-fold
OF YOUNG ADULTS WITH MAJOR DEPRESSION AND/OR ANXIETY HAVE MIGRAINE	OF PEOPLE WITH MIGRAINE HAVE A LIFETIME ANXIETY OR MOOD DISORDER HISTORY	INCREASE IN MAJOR-DEPRESSION RISK ASSOCIATED WITH MIGRAINE

These figures are not prevalence estimates for an individual clinic. Their value is to make comorbidity an active part of the routine review. Ask about low mood, loss of interest, guilt, hopelessness, anxiety, panic and self-harm separately from the pain history. Then ask how each condition affects the other: whether attacks precede withdrawal, whether anticipatory worry constricts activity between attacks, whether depression reduces self-care, and whether illness behaviour changes during periods of headache freedom. A longitudinal account is more informative than a single cross-sectional label.

There is an important nuance for examination answers and clinical reasoning. One account describes the depressive association as relatively specific to migraine rather than other severe headache disorders, while other material in the corpus contests that specificity. The reliable conclusion is the reciprocal association itself, not an overconfident claim that depression identifies migraine. A patient can have depression and another headache disorder; conversely, a patient with migraine may have normal mood between attacks. Comorbidity increases clinical vigilance but does not replace neurological diagnosis.

GOOD TO REMEMBER

When a patient with migraine reports “depression,” establish whether this denotes persistent syndrome-level mood change, attack-related dysphoria, avoidance driven by fear of the next attack, medication effects, or more than one of these. The management target becomes clearer only after that separation.

8.3 Anxiety, panic and the bipolar spectrum

The psychiatric association extends beyond depression. Migraineurs are more likely to have **generalised anxiety disorder, panic disorder and bipolar disorder**. This should prompt a proper psychiatric assessment, not a shortcut from headache to diagnosis. Anxiety may be an independent disorder, a response to an unpredictable and disabling illness, an affective component of an aura, or an acute reaction to a frightening neurological symptom. Panic

symptoms may similarly be reported around attacks without establishing that every episodic autonomic complaint is panic disorder.

The bipolar link needs especially disciplined language. Studies described in the source associate migraine specifically with phobias and with **hypomania/bipolar-spectrum conditions**, whereas the depression association may be found across headache disorders more generally. The evidence supplied here does not license a claim about a particular bipolar subtype. Do not convert "spectrum" into a more precise diagnosis than the history supports. Assess past periods of elevated or irritable mood, altered activity and their relation to headache independently, using the framework in the Mood disorders: classification, depressive and bipolar syndromes, neurobiology notes.

A useful formulation separates predisposing, precipitating and perpetuating influences without mistaking that structure for aetiological proof. Recurrent severe pain may increase vigilance to bodily sensations and reduce confidence in planning. Anxiety may then amplify avoidance, while low mood may shrink activity and social contact. These are clinically meaningful feedback loops even when no single symptom can be declared the original cause. The formulation should identify modifiable maintaining factors while retaining the neurological diagnosis as a real disorder in its own right.

For the same reason, avoid the outdated **"migraine personality"** formulation. Migraine is not confined to people of a particular social class or to those characterised as rigid, perfectionistic and competitive. Such a label can make the clinician mistake selection effects, coping style or current distress for a causal personality pattern. It also shifts attention away from the patient's actual syndrome. Personality formulation may still be clinically useful when it is based on the individual assessment, but it must not be smuggled in as an explanation for migraine.

TRAP **Do not diagnose bipolar disorder from migraine plus irritability alone.** Irritability can be part of an aura or an attack-state experience; a bipolar-spectrum diagnosis requires its own longitudinal evidence.

8.4 Aura and attack-state symptoms can resemble psychiatric phenomena

Migraine can transiently generate symptoms that fall naturally into psychiatric language. The psychiatric content of migraine aura includes anxiety, depression and irritability, alongside aphasia and altered perception of size, shape or time. During an attack, dysphoria and inattentiveness may look like a depressive or confusional state. Some patients withdraw; others appear unusually active. After recovery, an altered affective state may also be reported. The point is not to psychiatricise these symptoms, but to locate them accurately in the attack sequence.

Visual experience is a common source of confusion. **Visual migraine aura** may include a scotoma, flashing zigzag lines, coloured crescents, tubular vision or distortion of objects. A particularly useful discriminating feature is the **simultaneous positive and negative visual phenomena**: a scintillating experience may occur with an opaque or missing area. This pattern helps distinguish migraine aura from occipital seizure and from a psychiatric hallucination, but it

should be interpreted with the full time course and neurological context rather than used as a solitary test.

CLINICAL QUESTION	MIGRAINE-ORIENTED ANSWER	PSYCHIATRIC IMPLICATION
How does the symptom begin?	Ask whether it emerges gradually and evolves.	A gradual evolving symptom should prompt an aura history before a primary psychiatric inference.
What is experienced?	Look for positive phenomena alongside a deficit or obscuration.	Mixed positive and negative visual experience favours a migraine formulation in the stated differential.
What else is present?	Clarify sensory, language, perceptual and affective changes.	Fear, dysphoria or irritability may be attack-linked rather than independent syndromes.
What follows?	Establish the relationship to headache and recovery.	Repeated stereotyped sequencing is stronger evidence than a symptom label alone.

This approach is a diagnostic safeguard rather than an excuse to dismiss psychiatric symptoms. A patient may have an aura and a panic disorder, or migraine-related dysphoria and major depression. The relevant question is whether the symptom is confined to the migraine sequence, persists outside it, or has both forms. The Epilepsy and psychiatry notes should be used for seizure assessment; the Anxiety, OCD and trauma-related disorders notes cover diagnosis of the anxiety disorders themselves.

MNEMONIC

Sequence

For a possible aura, reconstruct the **sequence**: onset, evolution, paired positive and negative phenomena, headache relationship and recovery. The mnemonic is a history-taking spine, not a diagnostic criterion by itself.

8.5 Shared mechanisms: useful bridge, necessary caution

Mechanistic explanations can help the psychiatrist understand why a neurological pain disorder and psychiatric syndromes travel together, but they must not be represented as settled fact where the evidence only supports a proposal. In the model described by Kaufman, **spreading neuronal depression** begins with altered cortical neuronal metabolism and changes in neuronal activity that spread from posterior to anterior cortex. It is proposed to account for the aura and to activate the trigeminal system, followed by release of vasoactive neuropeptides and painful meningeal processes. This model displaced the older simple vascular account and places cortex, meninges and brainstem within the same clinical picture.

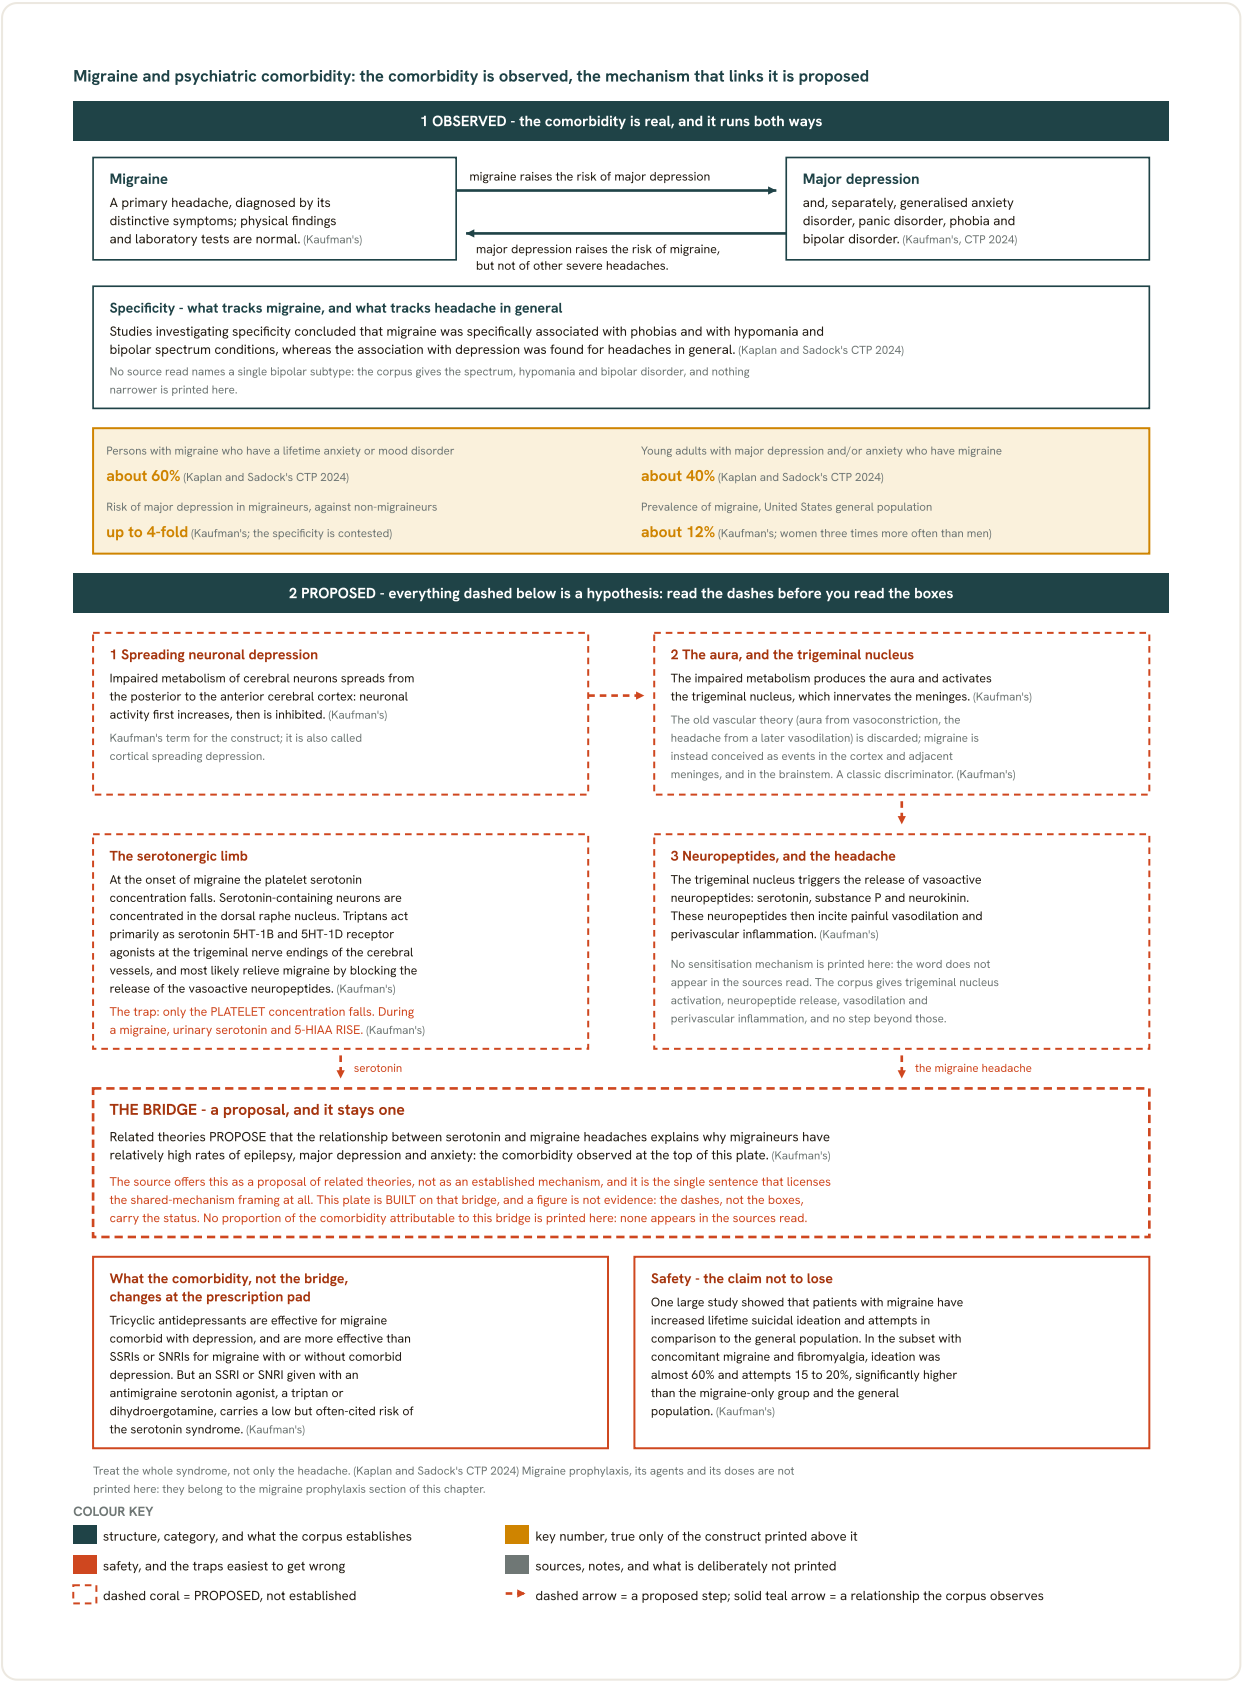


Fig 31.6 Migraine and psychiatric comorbidity: an observed comorbidity, and a bridge that is only proposed. Knowing which half is which is the discriminator, and the plate carries that on its face rather than in this caption: the solid teal half is observed, every dashed coral box and arrow is proposed. Three specifics are absent because they do not appear in the sources read, each marked where it belongs: any single bipolar subtype, any sensitisation mechanism, and any proportion of the comorbidity attributable to the bridge. (Kaufman's Clinical Neurology for Psychiatrists; Kaplan and Sadock's Comprehensive Textbook of Psychiatry 2024.)

The accompanying figure is best read as a conceptual bridge: a cortical event is linked to aura, trigeminal activation and pain, while shared biological systems may help explain why mood and anxiety disorders are frequent in the same patients. It is not a diagram of a proven unitary cause. In particular, it should not be used to infer that a patient's depression or anxiety is merely a migraine symptom.

The **serotonergic limb** adds a clinically familiar connection. Platelet serotonin concentration falls at migraine onset; serotonin-containing neurons are concentrated in the dorsal raphe nucleus; and triptans act primarily as 5HT-1B and 5HT-1D receptor agonists at trigeminal nerve endings, likely limiting the release of vasoactive neuropeptides. These observations make serotonin biologically plausible in migraine. They do not establish one shared abnormality as the cause of all comorbidity.

Related theories propose serotonin as a shared substrate for the association of migraine with epilepsy, major depression and anxiety. The word "propose" matters. This is a hypothesis that connects clinical observations and treatment pharmacology; it is not proof that serotonin explains the relationship in every patient. A good answer therefore gives the bridge, preserves the hedge, and returns to clinical formulation: migraine burden, psychiatric syndrome, temporal pattern, vulnerability and treatment exposure.

8.6 Risk assessment is part of migraine review

Comorbidity changes safety work. People with migraine have increased lifetime suicidal ideation and suicide attempts compared with the general population. In those with concomitant migraine and fibromyalgia, suicidal ideation is reported at **~60%** and suicide attempts at **15-20%**. The source also notes that the excess risk is not wholly explained by depression. This means that a clinician should not use the absence of an obvious major depressive syndrome as reassurance that suicide risk has been dealt with.

Ask directly about suicidal thoughts, intent, planning, past attempts, access to means and current supports whenever the clinical presentation indicates risk. Then integrate migraine into the formulation: severity and unpredictability of attacks, pain-related desperation, sleep disruption, functional losses, comorbid anxiety or mood symptoms, other pain syndromes and substance use. The headache history can clarify precipitants and burden, but it should never substitute for a full psychiatric risk assessment. Risk management follows the severity and immediacy of the person's presentation rather than the diagnostic label alone.

- Separate risk that intensifies during pain from risk that persists between attacks.
- Document the patient's account of functional losses and current supports.
- Reassess after changes in headache burden, psychiatric state or treatment exposure.

PITFALL

Do not reduce suicidal thoughts in a person with migraine to "understandable distress" and stop there. The association with suicidality requires direct assessment, including when depression is not the only apparent clinical problem.

8.7 Prescribing and diagnostic attribution at the interface

Medication choice should be made for the whole clinical syndrome, rather than by treating pain and psychiatric symptoms as unrelated lists. Tricyclic antidepressants are effective when migraine coexists with depression and are described as more effective for migraine than SSRIs or SNRIs. This is a comparative statement about migraine benefit, not a reason to ignore psychiatric indication, tolerability, previous response or the broader psychopharmacology plan. The Mood disorders: pharmacotherapy notes provide the general psychiatric pharmacology; migraine prophylaxis regimens are addressed separately in this chapter.

There is also a well-known interaction question. Concurrent use of an SSRI or SNRI with a triptan or dihydroergotamine carries a low but often-cited risk of **serotonin syndrome**. The clinically mature response is neither reflexive avoidance nor casual dismissal. Reconcile all medicines, identify the purpose of each serotonergic agent, explain relevant symptoms to the patient, and coordinate with the clinician treating the headache. Headache itself can also be an adverse effect of psychotropic medication, so a new headache after a psychiatric prescription deserves a temporal medication review rather than an automatic attribution to anxiety.

Prescribing at this interface is collaborative. The psychiatrist should make explicit which symptoms the psychiatric treatment is expected to target and which headache outcomes require review by the clinician managing migraine. This prevents one clinician from silently assuming the other has monitored the interaction burden, the patient's understanding or the effect of treatment on functioning. It also creates room to address psychological work: education about the distinction between attack symptoms and psychiatric symptoms, support for paced return to valued activity, and work with avoidance or catastrophic interpretations where these are present.

Diagnostic coding requires the same restraint. DSM-5 describes depression triggered by severe headaches as **Depressive Disorder due to Another Medical Condition**, but high migraine-depression comorbidity can make that causal claim difficult to defend. Use a "due to" formulation only when the evidence for causation is persuasive. Where reports of migraine are fabricated for opioid or another tangible benefit, the relevant alternatives named in the source are Factitious Disorder Imposed on Self and Malingering. These are serious determinations, not conclusions to draw from inconsistency or from a clinician's frustration with recurrent attendance.

> In migraine, psychiatric comorbidity is common and clinically consequential, but the diagnosis of a mood, anxiety or bipolar-spectrum disorder - and any causal attribution between it and headache - must rest on its own longitudinal evidence.

9 · Medication-overuse headache, tension-type headache and the psychiatric overlap

The psychiatric consultation can either interrupt or perpetuate a headache cycle. A patient may present with anxiety about a progressive head pain, low mood secondary to disability, insomnia, repeated requests for sedatives or analgesics, or apparent treatment resistance. The psychiatrist's task is not to annex the neurological diagnosis. It is to recognise when repeated symptomatic treatment is part of the disorder, distinguish a primary tension-type pattern from

transformed headache, and organise prescribing across settings so that one clinician does not unknowingly undo another's plan.

Chronic daily headache is a useful clinical syndrome rather than a shortcut to a single cause. It describes headache lasting 4 hours or longer on **at least 15 days each month** for at least 3 months. The central practical question is then retrospective: was this an episodic migraine or tension-type headache that gradually transformed, a newly persistent daily illness, a post-traumatic or psychiatric presentation, or a syndrome sustained by medicines? That history changes the plan more than the current pain adjective alone.

MOH	TTH	CBT
REPEATED RELIEF CAN SUSTAIN PAIN	COMMON PRIMARY HEADACHE	A TREATMENT PARTNER, NOT AN AFTERTHOUGHT

9.1 Medication overuse: a diagnostic habit of mind

Medication-overuse headache is best understood as a self-reinforcing clinical loop. It commonly develops on a pre-existing migraine substrate: acute medication is taken, rebound pain follows, and the rebound is interpreted as a reason for further self-administered medication. The patient is not necessarily seeking intoxication or behaving irrationally. They are pursuing immediate relief in a system in which immediate relief has become part of the longer-term problem. That formulation allows a firm discussion without moralising.

Several routes lead to chronic daily headache. Episodic migraine or tension-type headache may lose their episodic pattern; a new daily persistent presentation may begin in someone with little previous headache; post-traumatic and psychiatric contexts may contribute; and medication overuse can pave the route regardless of the original phenotype. In adults, prolonged exposure to headache medicines is the dominant risk factor. In children and adolescents, chronic daily headache usually occurs without such a history. This contrast prevents the psychiatrist from applying an adult explanation mechanically to a younger patient.

The medication history must include prescribed drugs, over-the-counter products, combination preparations, medicines borrowed from family, and tablets taken for sleep, distress or neck discomfort rather than explicitly labelled as headache treatment. The commonly implicated groups are **aspirin-butalbital-caffeine compounds**, triptans, non-steroidal anti-inflammatory drugs, benzodiazepines and opioids. This is psychiatry-facing neurology: butalbital-containing preparations and benzodiazepines may be introduced through psychotropic prescribing, while opioid exposure is often uncovered only when the history is asked in a non-accusatory way.

RULE **In a patient with increasingly frequent headache, every apparently helpful acute medicine is part of the diagnostic history.** Count the product, the indication the patient gives for it, the prescriber, and the sequence of relief and recurrence before adding another symptomatic agent.

Risk is not uniform across drug classes. Triptan or opioid use on **more than 10 days per month** is reported as a risk threshold for chronic daily headache, with development in less than 2 years. This is a risk statement, not a formal diagnostic criterion. By contrast, frequent non-opioid

analgesic use may require almost 5 years before chronic daily headache emerges. Barbiturate-containing products are particularly concerning: use on **at least 5 days per month** is associated with rebound headache and progression of headache disorders. Acute antimigraine medicines can also lead to chronic daily headache with use as little as 2 to 3 times a week.

PITFALL

Do not diagnose medication-overuse headache from a single threshold alone. The available frequency statements describe risk, not the formal per-class diagnostic criteria. Equally, do not dismiss a combination product because the patient calls it a "regular painkiller"; the sedative and caffeine components can obscure what is actually being overused.

9.2 Why the phenotype can deceive

Transformed chronic daily headache may look remarkably like tension-type headache at the consultation. Patients often describe generalised, waxing and waning, dull, pressing and non-pulsatile discomfort of mild-to-moderate intensity. Yet the discomfort can be incapacitating even when the pain is not dramatic. During transformation, the original headache may shed its distinctive qualities and associated aura-like or autonomic features. A cross-sectional history therefore risks labelling the current pain as uncomplicated tension-type headache and missing the antecedent migraine pattern and medication cycle.

Ask for the earliest memorable attacks, not only the current month: what did the first disabling headaches feel like, did activity worsen them, were there associated symptoms, and when did tablets become routine? Also ask what changed when the headache became almost daily: work attendance, family responses, sleep, food intake, and the patient's confidence in coping. These questions locate headache in its behavioural economy. They do not replace a neurological assessment, and the migraine disorder and its psychiatric comorbidity are addressed in the companion migraine discussion.

Psychiatric symptoms are neither incidental nor sufficient explanation. Chronic daily headache is associated with depression and anxiety, functional disability, stressful life events, disability, head trauma, obesity and socioeconomic adversity. In adolescents, depression and anxiety are the prominent risk factors. Depression is reported with chronic daily headache, migraine and tension-type headache, but not with trigeminal neuralgia or other secondary headache syndromes. A psychiatric diagnosis should therefore prompt a fuller headache formulation, not conclude it.

A useful formulation separates the pain disorder, the medicine-rebound loop, emotional distress and disability, and the care system that supplies medicines. Once separated, these processes can be addressed together without claiming that distress "caused" the pain. This is particularly important when a patient has been told that a bilateral pressing pain proves psychological tension. It does not.

9.3 Tension-type headache: recognise the pattern without psychologising it

Tension-type headache is the most common primary headache. It is often described as intermittent, bilateral frontal or cervical dull pain, with a pressing or squeezing quality. An episode may last 30 minutes to 7 days. The key clinical contrast with migraine is not simply severity. The patient typically has pain without the migrainous accompaniment of nausea, vomiting, photophobia, phonophobia, hyperacusis or autonomic disturbance.

"Tension" is historically sticky but clinically misleading. The disorder is not understood as the result of either muscle contraction or psychological tension. Neck discomfort may be a consequence of altered head movement in response to pain rather than proof that muscular tension caused it. This matters in psychiatry because a patient with stress and a headache can be given an etiological story that is both plausible and wrong. Stress may be clinically relevant, but relevance is not the same as a simple causal explanation.

MNEMONIC

PRESS

Pressing or tightening pain, **R**outine activity does not aggravate it, **E**mpty of prominent migrainous accompaniments, **S**ymmetric or bilateral pattern, **S**ustains usual activity. Use this as a bedside memory aid, then confirm the formal criteria and medication history.

Behaviour is a high-yield discriminator. People with tension-type headache usually tolerate the pain and continue daily activity; headache that is exacerbated by activity is more likely to be migraine. The disorders may nevertheless coexist, either together or in alternating periods. Many neurologists place them at opposite ends of a hypothesised spectrum arising from a shared but unknown physiological disturbance. That is a useful conceptual possibility, not an established mechanism and not permission to call every mild migraine "tension."

CLINICAL AXIS	TENSION-TYPE PATTERN	MIGRAINE-FACING CONTRAST
Location	Bilateral or symmetric pressure	May be hemicranial
Pain quality	Dull, pressing or tightening, non-pulsatile	Throbbing quality is more typical
Activity	Usually continues ordinary activity	Activity-related exacerbation favours migraine
Associated symptoms	Absence of prominent gastrointestinal and sensory concomitants	Nausea and sensory symptoms support migraine
Course	May coexist with migraine	Do not force an either-or history

9.4 Formal tension-type classifications

The **ICHD-3 frequent episodic tension-type headache** classification requires at least 10 episodes occurring on 1 or more but fewer than 15 days per month for at least 3 months. Episodes last from 30 minutes to 7 days. They have at least two of bilateral location, pressing or

tightening non-pulsatile quality, mild or moderate intensity, and no aggravation by routine activity. There is no nausea or vomiting, and no more than one of photophobia or phonophobia; another disorder must not better account for the headache.

ICHD-3 chronic tension-type headache occurs on **at least 15 days per month** for more than 3 months. The headache lasts for hours and may be continuous. The same core pain-character features apply. The chronic classification permits no more than one of photophobia, phonophobia or mild nausea, but excludes moderate or severe nausea and vomiting. Frequency alone cannot distinguish chronic tension-type headache from chronic migraine or chronic daily headache. Pain character, associated symptoms, prior phenotype and medicine use must do the discriminating work.

TRAP **A bilateral, pressing daily headache is not automatically chronic tension-type headache.** A transformed migraine with medication overuse can lose its original migrainous features and resemble the tension-type pattern. Ask backwards before labelling forwards.

The classification is a discipline of exclusion as well as description. It prevents the clinician from substituting a psychosocial impression for the actual symptom pattern. It also stops a false dichotomy: a person can have both migraine and tension-type headache, and an acute medicine used for either may contribute to chronic daily headache. A reliable formulation therefore names the current phenotype, the historical phenotype and the medication exposure separately.

9.5 Acute treatment and the prevention of iatrogenic chronicity

For tension-type headache occurring less than twice a week, acute treatment at onset with over-the-counter analgesics may suffice. Aspirin, aspirin-caffeine products, acetaminophen and non-steroidal anti-inflammatory drugs are named options. The apparently ordinary nature of this advice is precisely why it needs a ceiling in the patient's understanding: daily use often leads to chronic daily headache. The clinician should frame acute treatment as a rescue tool, not a standing response to every discomfort.

ACUTE OPTION	EVIDENCE STATEMENT WITHIN THIS SOURCE	MEDICATION-OVERUSE RELEVANCE
Ibuprofen 800 mg	First choice; lowest reported gastrointestinal bleeding or perforation risk among the compared options	Still requires a bounded acute-use plan
Naproxen sodium 825 mg	Next choice; higher gastrointestinal bleeding risk	Do not substitute repeated use for review
Aspirin or paracetamol 500 to 1000 mg	Comparable effectiveness in the cited comparison	Ask about all combined preparations
Caffeine 130 or 200 mg	Augments simple analgesics and ibuprofen in controlled trials	Combination products can conceal exposure

LOCUS is usually outpatient care, but its safety depends on a shared plan among all prescribers. FOCUS in the acute phase is to clarify what is being taken and stop automatic escalation; in the

short and middle term it is to restore a non-rebound pattern and treat disability, depression or anxiety; in the longer term it is to prevent recurrence of fragmented prescribing. MODUS includes medication review, psychologically informed education, behavioural treatment, and liaison with neurology and primary care. Relevant headache guidance should be consulted for an individual regimen, particularly when comorbidity or withdrawal risk is present.

GOOD TO REMEMBER

The most useful prescription may be a single reconciled medicine list. It should record every acute product, who supplied it, the symptom for which it is used, and the agreed clinician responsible for changing it. Fragmented prescribing is a mechanism of medication overuse, not merely an administrative inconvenience.

9.6 Withdrawal, containment and collaborative care

When chronic daily headache is attributed to medication overuse, the culprit medicine must be eliminated, but this is not a casual instruction to "just stop." Abrupt cessation can produce withdrawal symptoms as troublesome as opioid withdrawal. Rebound headaches may continue after detoxification, and relapse of medication use is common. The psychiatrist should provide containment and avoid reintroducing the same cycle simply to settle the immediate crisis.

LOCUS may need to shift from outpatient care to medically supervised withdrawal. Comorbid depression and a history of excessive use of any medicine are reasons neurologists often consider hospitalisation; withdrawal from over-the-counter preparations, opioids or excessive conventional antimigraine medicines may also require it. FOCUS is safety and containment in the acute phase, sustained discontinuation through the immediate rebound period, then relapse prevention and treatment of psychiatric comorbidity. MODUS is coordinated medical supervision, clear prescribing boundaries, psychological support, and a plan for the underlying headache disorder rather than a replacement dependence on another symptomatic medicine.

All of the patient's physicians should work together, actively seeking psychiatric disorder that may be obvious or initially hidden. This is not an instruction to psychologise the headache. Depression, anxiety, hopelessness about pain and medicine-related behaviours can make withdrawal harder, while uncontrolled pain and disability can intensify psychiatric symptoms. Chronic daily headache generally resists treatment; botulinum toxin has a role and tricyclic antidepressants, valproate and topiramate may help. Regimen selection belongs with the treating headache clinician. The psychotropic and anticonvulsant cross-use in migraine is covered separately in the migraine prophylaxis discussion.

9.7 Prevention in chronic tension-type headache

For chronic tension-type headache, preventative treatment may include amitriptyline or topiramate. When a tricyclic is selected for this indication, the cited approach is to start amitriptyline or clomipramine at 10 to 25 mg before bed and increase gradually. The reported average required amitriptyline dose for chronic tension-type headache is **50 to 75 mg per day**. These are tension-type headache prophylactic figures, not antidepressant dosing and not migraine-prophylaxis dosing.

The same source reports that most authorities recommend stopping treatment after 6 months regardless of efficacy, with a reduction of 20% to 25% every 2 to 3 days as a way to avoid rebound headache. LOCUS is outpatient care when risk is contained. FOCUS is prevention rather than repeated rescue medication; MODUS combines an individually prescribed preventive medicine, symptom monitoring, behavioural work and reconciliation of acute drugs. The general pharmacology of antidepressants and anticonvulsants belongs in the Mood disorders: pharmacotherapy notes and the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

Selective serotonin reuptake inhibitors have not been convincingly shown to prevent tension-type headache. This is a useful corrective to the tempting but simplistic idea that any antidepressant should help a disorder historically named for psychological tension. A tricyclic's benefit here should be discussed as a pain-preventive strategy, while depression and anxiety are assessed and treated on their own merits.

9.8 Behavioural treatment and the psychiatric contribution

Behavioural treatment is not a consolation prize for patients whose pain persists. Relaxation and biofeedback, alone or combined, have been tested in tension-type headache. Cognitive-behavioural stress-management approaches are effective, particularly for people with high stress when combined with relaxation or biofeedback. Improvement may appear more slowly than with pharmacological treatment, but can be sustained for several years without monthly therapist contact. Stress-management techniques combined with a tricyclic antidepressant are superior to either approach alone.

For children and adolescents with chronic headache, relaxation and cognitive-behavioural therapy reduce frequency and severity. In adults, insight-oriented psychotherapy and psychoanalysis do not alleviate the headache itself, although they may still provide insight, reduce anxiety, treat depression and offer other benefits. This boundary is clinically humane: do not promise headache relief from a therapy that is not shown to deliver it, and do not deny the therapy's other legitimate goals.

- Validate pain as pain, even when stress worsens coping or disability.
- Make acute-medication use visible before attempting to change it.
- Formulate depression and anxiety as comorbidity that can amplify suffering and complicate withdrawal.
- Use behavioural strategies to expand coping and reduce reliance on repeated rescue medication.

A daily bilateral pressing headache should trigger a retrospective migraine history and a complete medicine count before it is called chronic tension-type headache.

10 · Psychotropics in migraine prophylaxis and antidepressant or anticonvulsant cross-use

The examination issue is not simply which psychotropic can prevent migraine. The higher-value task is to keep indication, dose and therapeutic purpose aligned when the same medicine crosses neurology and psychiatry. Amitriptyline may be chosen for migraine in a person who is depressed, and topiramate may be familiar from anticonvulsant practice, but neither familiarity licenses transfer of a dose from one indication to another. This is prescribing by indication, not prescribing by drug name.

The psychiatrist therefore needs a compact prophylaxis framework: identify the preventive indication, select a suitable agent after discussing benefit and risk, write the regimen with its migraine scope, and review safety in the context of mood symptoms, self-harm risk, pregnancy potential and concurrent prescribing. The disorder, its psychiatric comorbidity and its shared mechanisms are considered elsewhere in this chapter. General psychopharmacology belongs in the Mood disorders: pharmacotherapy notes, while the general beta-blocker and anxiolytic profile of propranolol belongs in the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes.

RULE **A cross-use drug never carries a context-free dose.** Write the indication in the same sentence as the dose, because a familiar drug name can conceal a materially different therapeutic range and purpose.

10.1 Read the prescription as an indication-dose pair

A psychotropic used in migraine prophylaxis sits at an interface between specialties. The prescription may look familiar to a psychiatrist, yet the clinical question has changed: the target is prevention of migraine attacks rather than treatment of a mood syndrome or seizure disorder. This changes how the dose is interpreted, how response is judged and which safety discussion has priority. A correct answer therefore begins with the indication and only then names the agent and regimen.

NICE migraine-prevention options are propranolol, topiramate or amitriptyline after a full discussion of the benefits, risks and suitability of each. This is a choice among reasonable options, not a command to use every option in a fixed sequence. The discussion should incorporate the patient's previous treatment experience, concurrent medicines, likely acceptability and risks that become especially important in psychiatric practice.

Cross-use can be clinically useful because one prescriber may recognize more than one relevant domain. It can also create anchoring: seeing amitriptyline and assuming an antidepressant regimen, or seeing topiramate and assuming an anticonvulsant regimen. The disciplined response is to ask, “What is this medicine being used to prevent or treat here?” before reading any number.

MNEMONIC

IDI: Indication, Dose, Interpretation

Indication tells you which regimen is valid; **D**ose must be written with that scope; **I**nterpretation asks whether the observed benefit belongs to migraine prevention rather than to another property of the medicine.

In migraine prophylaxis, never import the psychiatric or anticonvulsant dose merely because the drug is shared across specialties.

10.2 Propranolol: the reference preventive and the overdose problem

Propranolol is the most commonly used migraine prophylactic in *Essentials of Medical Pharmacology*. For *migraine prophylaxis*, it is started at **40 mg twice daily** and may be increased up to **160 mg twice daily**. Those figures belong to migraine prevention. They do not describe its anxiolytic use, and this fragment does not reproduce its general beta-blocker pharmacology.

The reason propranolol provides a useful reference point is that it gives the candidate a benchmark against which the cross-use agents can be understood. Its response figures can be set beside the attack-reduction evidence for topiramate, as in the strip below. The comparison is about migraine outcomes, not about which drug is more familiar to the psychiatrist. Response assessment should therefore remain tied to the preventive target rather than to nonspecific impressions of feeling better or worse.

UP TO 70%	4 WEEKS	50% REDUCTION
PROPRANOLOL: FREQUENCY AND SEVERITY	PROPRANOLOL: EFFECT GENERALLY SEEN	TOPIRAMATE: HALF OF PATIENTS ACROSS TWO RANDOMISED TRIALS

The central psychiatric complication is not a subtle pharmacological distinction. In a person with depression and migraine, **propranolol overdose risk** may increase the danger of self-harm because toxicity and deterioration after overdose can be rapid. NICE advises caution and harm-minimising prescribing in precisely this comorbid group. A preventive choice that is reasonable neurologically may therefore be unsuitable when access to medication, current suicidal intent, impulsivity or the feasibility of safe dispensing changes the risk calculation.

GOOD TO REMEMBER

Before prescribing propranolol for migraine to a depressed patient, make overdose safety part of drug selection rather than an afterthought. The practical response may include choosing a safer option, limiting the amount accessible at one time and involving supports where clinically appropriate. These are risk-management principles, not substitutes for a full suicide-risk assessment.

10.3 Amitriptyline: migraine benefit is not antidepressant treatment

For *migraine prophylaxis*, amitriptyline is prescribed at **25 to 50 mg at bedtime**. The scope is essential: this is not an antidepressant-dose statement. General pharmacology and psychiatric

dosing should be read in the Mood disorders: pharmacotherapy notes. Importing a mood-disorder regimen into a migraine answer is not an ambitious answer; it is a category error.

The decisive concept is that the **antimigraine effect of amitriptyline** is independent of its antidepressant property. A patient does not need to improve from depression for the preventive migraine effect to be clinically meaningful, and a migraine response does not demonstrate that depression has been adequately treated. The two outcomes require separate assessment even when one medicine appears on both problem lists.

Amitriptyline can still be attractive when migraine and depression coexist because the class suits that clinical combination. “Suits” does not mean that one migraine regimen automatically treats both conditions. It means the prescriber may consider the overall fit while keeping indications, outcome measures and dose decisions separate. It is effective in many patients but causes **more adverse effects than propranolol**, so apparent convenience must be weighed against tolerability.

TRAP **“Amitriptyline prevents migraine because it treats the associated depression.”** The grounded discriminator is the opposite: its antimigraine action is independent of its antidepressant property. Comorbidity may influence selection, but it does not explain the preventive effect.

This distinction also prevents false conclusions during follow-up. If headache burden improves while depressive symptoms persist, the migraine regimen may be working and the mood disorder may still require its own treatment plan. If mood improves while migraine does not, the shared drug name should not obscure failure of the preventive target. One medicine can therefore generate two clinically separate response narratives.

10.4 Anticonvulsant cross-use: sequence by indication, not familiarity

Valproic acid, gabapentin and topiramate illustrate anticonvulsant cross-use. In migraine, the cited pharmacology source places anticonvulsants below beta blockers in efficacy and uses them particularly when other drugs have not worked or propranolol is unsuitable. That **lower efficacy than beta blockers** is a sequencing statement for migraine prevention; it is not a judgment about their value in seizure treatment.

For *migraine prophylaxis*, valproic acid is used at 400 to 1200 mg per day. For the same indication, gabapentin has some preventive effect at 300 to 1200 mg per day. These doses should not be expanded into general pharmacology, monitoring or psychiatric prescribing rules here. Valproate's broader use belongs in the Mood disorders: pharmacotherapy notes, and anticonvulsant treatment more generally belongs in the Epilepsy and psychiatry notes.

Topiramate makes the dose-scope problem especially visible. For *migraine prophylaxis*, **topiramate titration** begins at **25 mg once daily** and is gradually increased to **50 mg once or twice daily**. By contrast, for the *anticonvulsant indication*, the dose may be increased up to 100 to 200 mg twice daily. The latter can be four times the migraine-prevention target and must never be printed as though it were the migraine regimen.

Evidence of preventive benefit should also be read in migraine terms. The trial result in the evidence strip communicates the kind of outcome being pursued: fewer attacks, not seizure control and not a psychiatric effect. It should not be converted into a claim about anticonvulsant efficacy or psychiatric benefit. Topiramate's broader pharmacology, including its relevance to weight, cognition and raised intracranial pressure, is considered in the idiopathic intracranial hypertension section of this chapter.

PITFALL

A medication reconciliation that records only "topiramate" without the indication invites a dosing error. Record migraine prophylaxis explicitly, verify why any concurrent prescriber is using it, and reconcile the intended target before altering the regimen.

10.5 Flunarizine and the Indian-practice divergence

Guideline lists are shaped by jurisdiction and should not erase familiar local practice. NICE names propranolol, topiramate and amitriptyline for consideration, but it does not name flunarizine. Indian pharmacology teaching and practice nevertheless include **flunarizine for migraine prophylaxis** as a standard option. The correct synthesis is to present the UK recommendation accurately and then identify the Indian option separately, without pretending that both documents contain the same list.

IN INDIA

For *migraine prophylaxis*, flunarizine is used at 10 to 20 mg once daily in adults and 5 mg once daily in children. Its routine place in Indian practice is a genuine jurisdictional difference from NICE, not an omission to be silently "corrected" in either direction.

This comparison teaches a broader method for postgraduate answers. State the named guideline's recommendation as that guideline's recommendation. Then state the locally relevant alternative and its source. A candidate loses precision by merging the two into an imaginary universal ladder. A clinician loses more: a locally available and familiar option may be missed, or a foreign recommendation may be applied without considering access and context.

The same discipline applies to children. A paediatric dose grounded for flunarizine does not authorise extrapolation of adult regimens for other medicines. In June 2025, the use of **topiramate and amitriptyline for migraine prevention in children and young people was off-label** under the cited NICE guidance. Off-label status calls for explicit prescribing responsibility and explanation; it is not synonymous with ineffectiveness or prohibition.

10.6 Pregnancy potential and self-harm risk are selection variables

Preventive prescribing is not complete when the drug and dose are written correctly. Psychiatric practice changes the salience of two safety questions: could the available medicine be used in self-harm, and could the regimen expose a pregnancy? These questions belong at the point of selection because discovering them after dispensing may leave the patient with a theoretically suitable but practically unsafe option.

For topiramate, the restriction is categorical in the cited guidance. It **must not be used for migraine prophylaxis during pregnancy**. In women of childbearing potential, it must not be used unless the conditions of the **Pregnancy Prevention Programme** are fulfilled, including highly effective contraception. This is especially important when the psychiatrist encounters topiramate on a medication list and might otherwise regard it as merely a familiar cross-specialty drug.

For propranolol, the psychiatric safety lens is different. The concern is access to a medicine capable of serious toxicity in overdose in a population that may have depression and self-harm risk. The response is individualised risk minimisation rather than automatic exclusion of every depressed patient. Ask about current and past self-harm, intent, impulsivity, available quantities and who controls medication access, then incorporate that information into the preventive choice.

- **Before selection:** establish the migraine-prevention target, current psychiatric state, pregnancy potential and overdose context.
- **At prescribing:** document the indication beside the regimen and explain why this option fits the individual risk profile.
- **At review:** assess migraine outcome separately from mood, seizure or nonspecific symptom change.
- **Across prescribers:** reconcile who owns titration, monitoring and renewal so that a dose from another indication is not imported inadvertently.

10.7 Comparative prescribing map

The doses have been taught at their primary clinical locations above. The table instead compares the reason each agent appears in a psychiatry-facing migraine discussion, the boundary that prevents indication transfer, and the application most likely to alter a prescribing decision. It is not a complete pharmacology table and should not be used to infer unlisted contraindications, adverse effects or monitoring schedules.

AGENT	PLACE IN MIGRAINE PREVENTION	CROSS-USE BOUNDARY	PSYCHIATRIC-INTERFACE APPLICATION
Propranolol	Most commonly used preventive in the cited pharmacology source	Its migraine regimen must not be confused with anxiolytic use	Overdose toxicity and rapid deterioration matter when depression or self-harm risk is present
Amitriptyline	Effective in many patients, but with more adverse effects than propranolol	Its preventive effect is independent of antidepressant action	Assess migraine response and depressive symptoms as separate outcomes
Flunarizine	Standard Indian option absent from the NICE list	This is a jurisdictional difference, not antidepressant or anticonvulsant cross-use	State the Indian option separately from UK guidance
Valproic acid	Anticonvulsant option within a class whose migraine efficacy is lower than beta blockers	Its mood-disorder and anticonvulsant uses do not define the migraine regimen	Keep general pharmacology in its owning chapters
Gabapentin	Has some migraine-prophylactic effect	Familiarity from anticonvulsant practice does not make indications interchangeable	Judge benefit against the migraine target
Topiramate	Preventive option with attack-reduction evidence	The migraine target must remain distinct from the higher anticonvulsant dose	Apply pregnancy and Pregnancy Prevention Programme restrictions
Erenumab	Migraine-specific preventive rather than psychotropic cross-use	It illustrates a different type of preventive decision	Indian availability and cost require local verification

Erenumab is included for completeness because it represents a migraine-specific preventive rather than psychiatric or anticonvulsant cross-use. It is a fully human monoclonal antibody that blocks the **CGRP receptor**. For *migraine prophylaxis*, it is given as 70 mg subcutaneously once a month and significantly reduced monthly migraine days in people with 4 to 14 migraine days per month. Shared drugs require unusually careful dose scoping; this migraine-specific agent instead raises practical questions about local access. Indian availability and affordability are not established by the allowed evidence and should be checked rather than assumed.

10.8 Management: LOCUS, FOCUS and MODUS

Management begins in the outpatient LOCUS. Migraine prophylaxis is usually a planned, longitudinal decision, while emergency or inpatient care becomes relevant when psychiatric risk, overdose or another acute safety issue changes the setting required. The guideline anchor for the choice is NICE CG150, which recommends considering propranolol, topiramate or amitriptyline

after discussing benefits, risks and suitability. Indian practice must additionally retain flunarizine as a locally standard option rather than treating the NICE list as globally exhaustive.

FOCUS. The immediate focus is safe selection and an unambiguous indication-dose prescription. The short-term focus is tolerability, adherence and whether the preventive plan is being followed as intended. The continuing focus is reduction in migraine burden while separately tracking psychiatric symptoms, pregnancy-related safety and self-harm risk. The purpose of separation is not bureaucratic: it prevents a response in one domain from being mistaken for adequate treatment in another.

MODUS. Drug treatment is the owned mode in this fragment. Psychological and social care remain important when stress, mood symptoms, adherence, family support or self-harm risk affect the treatment plan, but this section does not duplicate psychotherapy or full migraine management. Procedural options are not specified in the allowed evidence. Collaborative care means the psychiatrist, physician or neurologist agrees on the indication, chosen regimen, titration responsibility and review target.

The prescribing sequence is simple. Write “for migraine prophylaxis” beside every shared-drug regimen; choose among appropriate preventives only after discussing individual benefit, risk and suitability; screen specifically for propranolol access in a self-harm context and for topiramate restrictions related to pregnancy potential; do not treat comorbidity as proof that one regimen is adequately treating both conditions; and cross-refer general psychopharmacology rather than importing its doses into the migraine plan.

The final examination answer can be concise without being thin: state the preventive choices, give only indication-scoped regimens, identify amitriptyline's independent antimigraine effect, contrast the migraine and anticonvulsant topiramate doses, name flunarizine's Indian place, and close with the propranolol overdose and topiramate pregnancy safeguards. That sequence demonstrates clinical reasoning rather than a memorised list.

11 · Idiopathic intracranial hypertension and psychotropic associations (topiramate, weight)

The marks lie at the interface between headache, threatened vision, obesity and psychopharmacology. **Idiopathic intracranial hypertension** is the preferred name for raised intracranial pressure without an identified cause. The older labels, pseudotumor cerebri and benign intracranial hypertension, are historically recognisable but clinically misleading: there is no tumour, yet the disorder is not benign because untreated optic-disc swelling can culminate in irreversible loss of vision. For the psychiatrist, the practical task is not to manage every headache as neurological disease. It is to recognise the pattern that warrants examination of the optic discs, protect vision while the diagnosis is established, and review medication and weight without making a premature causal attribution.

This topic also corrects an attractive but unsafe shortcut. Obesity is strongly associated with the disorder, and topiramate can suppress appetite and reduce weight. That does not establish

topiramate as treatment for idiopathic intracranial hypertension. The verified treatment statement names acetazolamide as first-line. The standard topiramate reference does not recommend combining topiramate with acetazolamide. The useful lesson is therefore a prescribing tension, not a therapeutic equivalence.

RULE **In headache with visual symptoms, examine or arrange examination of the optic discs before debating psychotropic causation.** A normal-looking general neurological examination does not neutralise papilloedema or its threat to sight.

11.1 Clinical phenotype and why "benign" is wrong

The typical clinical setting is a young woman with obesity, sometimes with menstrual irregularity. Female sex and obesity carry the most robust association. This demographic pattern should raise suspicion, but it must not become a stereotype: the diagnosis rests on the clinical syndrome, examination, imaging, cerebrospinal-fluid pressure and exclusion of another cause. Conversely, psychiatric comorbidity or psychotropic exposure must not be allowed to make a neurological complaint seem functional by default.

The characteristic symptom cluster is a dull, generalised headache, pulsatile tinnitus and transient episodes in which vision briefly darkens. These visual obscurations are especially important because patients may describe them vaguely as "blackouts", "giddiness" or anxiety. Careful history distinguishes a transient visual phenomenon from loss of consciousness. Nausea or vomiting may join the syndrome. Persistent evolution matters more than a single dramatic adjective: a headache that persists and accumulates visual or gastrointestinal features deserves a different response from a familiar episodic migraine-like pattern.

CLINICAL DOMAIN	GROUNDED FINDING	PSYCHIATRIC IMPLICATION
Optic discs	**Papilloedema** is the cardinal sign	Fundoscopy cannot be replaced by a reassuring mental-state or general examination
Symptoms	**Dull generalised headache, pulsatile tinnitus and transient visual obscurations**	Ask what "blackout" means before classifying it as panic, dissociation or syncope
Neurological examination	It may otherwise remain normal despite florid optic-disc swelling	Normal power, reflexes, coordination and gait do not close the case
Visual consequence	An enlarged blind spot may progress to optic atrophy and blindness if untreated	The referral endpoint is preservation of sight, not merely relief of headache

Raised pressure may stretch the **abducens nerve and produce unilateral or bilateral inward deviation of the eye**. This is a classic false-localising sign: it is compatible with the syndrome and does not, by itself, prove a focal pontine lesion. Its presence may generate diplopia and makes the eye examination still more important. The broader framework for raised intracranial pressure and hydrocephalus belongs in the Stroke, TBI, delirium and focal brain syndromes notes.

PITFALL

Do not let “neurological examination otherwise normal” become “no neurological disease”. The decisive abnormality may be on optic-disc examination, while the rest of the bedside examination remains unexpectedly unremarkable.

11.2 Diagnosis is a structured exclusion

The **Modified Dandy criteria** organise the diagnosis around **six** elements: symptoms or signs of raised pressure; no localising sign apart from abducens palsy; an awake and alert patient; imaging without an explanatory lesion or thrombosis; raised lumbar-puncture opening pressure with normal cerebrospinal-fluid composition; and no alternative explanation. The diagnostic opening-pressure threshold is greater than **25 cm H2O**. This is a criterion, not a licence to diagnose from pressure alone. The clinical context, imaging and fluid composition remain part of the same diagnostic statement.

MRI with MR venography is used to exclude an intracranial mass and interrogate venous drainage. Imaging may show compressed, slit-like ventricles, while venography may show slowing or occlusion of intracranial venous drainage. These findings do not convert every case into idiopathic disease. “Idiopathic” is a residual label: when a venous sinus thrombosis or another cause explains the pressure, the case is reclassified as secondary intracranial hypertension. Detailed modality selection and sequence interpretation are covered in the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

The lumbar puncture contributes pressure measurement and analysis of the fluid. In idiopathic intracranial hypertension the glucose remains normal and cells are absent; protein may be low. Thus the defining abnormality is pressure in a fluid that is otherwise essentially normal. A cellular, biochemical or other unexpected abnormality should prompt reconsideration rather than being forced into the idiopathic label. Cerebrospinal-fluid technique and the general neuropsychiatric work-up are addressed elsewhere in this chapter.

MNEMONIC

PIVOT: Persistent headache, Inspect discs, Venous imaging, Opening pressure, Threat to sight

Persistent evolving headache with visual symptoms; **I**nspect the optic discs; use **V**enous imaging while excluding another cause; interpret the **O**pening pressure within the full diagnostic criteria; remember the **T**hreat to sight.

11.3 Management: sight is the endpoint

LOCUS. A stable patient without rapidly progressive visual concern may enter care through outpatient neurology or ophthalmology, but new visual loss, significant optic-disc swelling or clinical deterioration requires urgent specialist assessment. The psychiatrist’s locus is recognition, medication reconciliation, communication and continuity of mood treatment. Definitive

diagnostic and vision-preserving decisions belong with clinicians who can examine the discs, assess visual function, arrange imaging and perform lumbar puncture.

FOCUS. The acute focus is to confirm raised pressure, exclude an identifiable cause and protect sight. The short-term focus is control of the pressure syndrome and objective surveillance of visual risk. The continuing focus is to prevent optic atrophy, review relevant comorbidity and medication exposure, and avoid destabilising psychiatric illness through reflex prescribing changes. Headache relief matters, but it is not the sole or safest measure of success.

MODUS. The grounded pharmacological approach uses carbonic-anhydrase inhibitors and diuretics, with **acetazolamide as first-line treatment**. Repeated lumbar puncture may be used by neurological teams to measure pressure and drain cerebrospinal fluid. Refractory disease may require a cerebrospinal-fluid shunt or optic-nerve-sheath fenestration; the procedural purpose is prevention or slowing of visual loss. Weight and medication review are relevant supportive actions, but they do not substitute for pressure-directed and sight-directed care.

- Refer for a clinician who can assess papilloedema and complete the diagnostic lumbar-puncture evaluation.
- Review all medicines and comorbidities before assigning causality to a psychotropic.
- Coordinate any mood-stabiliser change with the patient and the neurological team rather than acting on association alone.

GOOD TO REMEMBER

A patient may report less headache while vision remains at risk. Keep the management target anchored to the optic nerves and visual function, especially when procedural escalation is being considered.

11.4 Lithium: association is not simple causation

Lithium is among the medication exposures associated with idiopathic intracranial hypertension, alongside tetracycline-related antibiotics, hormonal contraceptives and vitamin A. Association, however, is not proof that lithium caused an individual case. This distinction is crucial because people receiving lithium may also have obesity, hypertension, chronic kidney disease or another relevant exposure. Some published cases also involved another implicated medicine, and reported lithium concentrations could remain therapeutic.

0.03 to 4.7 per 100,000	almost exactly 9:1	7.70 → 4.75	17
COUNTRY-SPECIFIC ADULT INCIDENCE	FEMALE:MALE RATIO IN A SWEDISH DATABASE PERIOD	LITHIUM ODDS RATIO: GENERAL TO OBESE CONTROLS	LITHIUM-ASSOCIATED CASES IN A REVIEW

The control-group comparison is the teaching point. Against obese controls rather than general-population controls, the odds ratio falls from 7.70 to 4.75, which is a reduction of roughly two fifths of the apparent lithium signal. That pattern shows how strongly obesity can confound an observed drug association. Residual confounding by hypertension and chronic kidney disease

further limits causal certainty. The small published case count, co-prescribed implicated drugs and therapeutic lithium concentrations all argue against inflating a rare association into an expected toxicity.

Routine screening for idiopathic intracranial hypertension is not required simply because lithium is prescribed. The clinically useful posture is selective suspicion. Ask specifically about a persistent headache that is new or evolving, rather than assuming every episodic migraine-like headache represents raised pressure. Escalation is warranted when persistence is joined by nausea, vomiting or transient visual change. Referral is needed for papilloedema assessment and lumbar puncture when the syndrome raises concern.

TRAP **Do not make lithium discontinuation the first response to a diagnosis of idiopathic intracranial hypertension.** First examine competing medicines, obesity and relevant comorbidity, then discuss causality and viable mood-stabiliser alternatives with the patient and neurologist.

11.5 Weight is the shared clinical hinge

The overlap between psychiatry and idiopathic intracranial hypertension is partly created by weight. Obesity is among the strongest associated factors for the neurological disorder. A weight-gain problem arising during psychotropic treatment can potentially strengthen that background risk, while topiramate may suppress appetite and produce weight loss. The metabolic mechanisms and management of psychotropic-associated weight gain belong in the Endocrine and metabolic psychiatry notes; the neurological task here is to recognise why weight trajectory changes the risk conversation.

Topiramate-associated appetite suppression and weight loss can be desirable in a patient who is overweight or obese, but problematic in someone who is underweight. In epilepsy studies, the weight effect was more evident at higher baseline weight and was greater in females than in males of similar baseline weight. That is strikingly aligned with the population in whom idiopathic intracranial hypertension is most often encountered. It makes topiramate clinically interesting, but does not turn a shared demographic pattern into evidence of disease treatment.

Controlled studies also support a weight effect in psychiatric populations, including reduction of adiposity in women receiving olanzapine. This can inform a broader metabolic plan, but it must remain properly scoped. A medicine selected to address weight or another established indication is not thereby a pressure-lowering treatment. No topiramate regimen for idiopathic intracranial hypertension is established by the source base used here, and migraine-prophylaxis dosing belongs to the dedicated headache pharmacotherapy discussion.

RULE **Weight benefit and disease-specific efficacy are different constructs.** Topiramate may help weight in the same demographic that is vulnerable to idiopathic intracranial hypertension. That overlap cannot be used to claim that it treats the neurological disorder.

11.6 Why topiramate is not acetazolamide

Topiramate has several pharmacodynamic actions: state-dependent sodium-channel blockade, positive modulation of GABA-A receptors, and antagonism at kainate and AMPA receptors, together with **weak inhibition of the CA-II and CA-IV carbonic-anhydrase isoenzymes**. The relation between these biochemical actions and each clinical effect is not fully established. The word "weak" is load-bearing. It prevents a category error in which any carbonic-anhydrase inhibition is treated as equivalent to acetazolamide.

The verified consequences attached to topiramate's carbonic-anhydrase effect are not a cerebrospinal-fluid mechanism. They include reduced urinary citrate, increased urinary pH and a contribution to nephrolithiasis; the same action may contribute to paraesthesia. About **1.5%** of exposed adults develop nephrolithiasis, and adequate fluid intake may reduce risk. Carbonic-anhydrase inhibition also underlies concern about hyperchloraemic non-anion-gap metabolic acidosis. Baseline and periodic serum bicarbonate measurement is recommended, with dose reduction, discontinuation or alkali considered when acidosis is detected.

Acetazolamide shares clinically recognisable adverse effects such as paraesthesias and kidney stones, and may also cause gastrointestinal upset, hypokalaemia, rash and, rarely, aplastic anaemia. This overlap helps explain why stacking carbonic-anhydrase inhibitors is not a neutral act. In a patient whose neurological treatment already includes acetazolamide, adding topiramate purely because weight loss seems attractive creates avoidable interaction risk.

PITFALL

Combining topiramate with acetazolamide is not recommended because the combination can increase the risks of metabolic acidosis, nephrolithiasis and heat-related problems. Do not convert "each inhibits carbonic anhydrase" into "combine them for greater benefit".

11.7 Topiramate adverse effects can imitate the problem

The **topiramate cognitive and neuropsychiatric adverse-effect signature** includes confusion, psychomotor slowing, impaired concentration or attention, memory difficulty, speech or language problems, depression or other mood problems, sedation and fatigue. Patients often experience this as slowed thinking or difficulty finding words rather than volunteering the label "cognitive adverse effect". These costs matter in postgraduate trainees, professionals, students and anyone whose recovery depends on cognitive performance. They must be weighed against any weight benefit.

Topiramate is also associated with increased risk of suicidal ideation or behaviour across epilepsy and psychiatric populations. Assessment should therefore include change in mood, hopelessness and suicidality, not only weight and headache. Paraesthesia and renal stones may arise from the same carbonic-anhydrase pharmacology discussed above. A new somatic complaint after starting the medicine should be characterised rather than automatically attributed to anxiety or to the intracranial-pressure disorder.

The most important visual discriminator is drug-induced ocular disease. Topiramate is rarely associated with **acute myopia and secondary angle-closure glaucoma**, presenting with acute reduction in visual acuity or ocular pain, usually within **one month** of initiation and reversing with discontinuation. Idiopathic intracranial hypertension, by contrast, threatens vision through papilloedema, visual-field change and eventual optic atrophy. A patient taking topiramate can therefore develop visual symptoms through a mechanism that is not idiopathic intracranial hypertension.

- Transient visual obscurations with persistent headache direct attention toward papilloedema and raised pressure.
- Acute ocular pain or abrupt loss of acuity after starting topiramate raises concern for its ocular adverse reaction.
- Either presentation warrants prompt eye assessment; neither should be dismissed as a nonspecific medication complaint.

11.8 The lithium, acetazolamide and topiramate junction

A lithium-treated patient who develops idiopathic intracranial hypertension sits at a difficult prescribing junction. Lithium is an associated exposure, although the evidence is tenuous; acetazolamide is first-line neurological treatment; and topiramate may be considered elsewhere in the psychiatric plan because of weight. Each decision changes the interpretation or monitoring of the others. The clean approach is to separate diagnosis, causality, neurological treatment and metabolic strategy rather than trying to solve all of them with a single drug switch.

The acetazolamide interaction evidence comes from oedema management rather than an idiopathic-intracranial-hypertension treatment study: acetazolamide lowers lithium concentrations by **31%**. The practical inference is that a lithium-treated patient started on acetazolamide requires serum-level monitoring and dose adjustment guided by measured concentrations. Topiramate, particularly at high dose, may also affect serum lithium concentrations; its prescribing information recommends monitoring lithium when the medicines are co-administered in that setting. The clinical importance of this interaction remains incompletely characterised, so certainty should not be overstated.

A disciplined review separates the questions that are often collapsed. First, confirm that the raised-pressure diagnosis is secure and that an identifiable secondary cause has been excluded. Next, ask which medicines and comorbidities plausibly contribute, and grade the strength of evidence for each rather than treating every association equally. Finally, establish what the neurological team needs to protect vision and what drug-level or biochemical monitoring follows from that treatment. This sequence preserves diagnostic clarity while allowing urgent care to proceed.

The psychiatrist then protects continuity of effective mood treatment while remaining willing to change it when the individual causal assessment supports doing so. Obesity, hypertension, chronic kidney disease, other implicated medicines and viable alternatives all belong in the shared decision. The patient should hear a balanced explanation: lithium has been associated

with the disorder, but the association is uncommon and confounded; stopping it reflexively is not the recommended first move; and vision-preserving neurological care must proceed without waiting for perfect causal certainty.

> Memorise the tension: acetazolamide is first-line for idiopathic intracranial hypertension; topiramate is only a weak carbonic-anhydrase inhibitor. It is not established here as treatment for the disorder and should not be combined with acetazolamide.

Sleep: the architecture and its disorders

Sleep neurophysiology, breathing, circadian rhythm and investigations

Why these marks belong in neurology. Sleep is not diagnosed from stillness alone. It is a changing physiological state expressed through cerebral activity, eye movements, skeletal-muscle tone, ventilation, oxygenation and observable behaviour. The neurological task is to identify the system that controls state, recognise when breathing or timing is disturbing that state, and choose a recording that can make the suspected disturbance measurable. This creates a continuous clinical argument: circuit, clock and breathing determine what happens during sleep; the investigation converts each of those processes into a signal that can be examined.

This section deliberately provides that neurological frame rather than duplicating a sleep-medicine chapter. Detailed sleep architecture, the hypnogram, formal staging, the two-process model, sleep-disorder criteria and treatment algorithms belong in the Eating and sleep-wake disorders notes. Molecular circadian physiology and melatonin dynamics belong in the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes. Here, the emphasis is on the landmarks a psychiatrist should recognise and the reasoning that links a clinical question to a physiological channel.

STATE	BREATH	CLOCK	TRACE
AROUSAL CIRCUIT	VENTILATORY INTEGRITY	CIRCADIAN ALIGNMENT	PHYSIOLOGICAL CORRELATE

12.1 State control: the neurological map

The useful starting point is not a list of sleep stages but a map of competing state-control systems. The **transmitter pattern links dopamine, serotonin, histamine and acetylcholine with wakefulness; reduced dopamine and serotonin with increased GABA activity characterises both NREM and REM sleep, while acetylcholine remains active in wakefulness and REM rather than NREM**. This is a state association, not a licence to infer sleep from the concentration of any single transmitter. It tells the candidate why sleep and wakefulness are network phenomena and why a drug acting at one transmitter may alter vigilance without reproducing normal sleep physiology.

Hypocretin, also called orexin, maintains wakefulness rather than initiating it. Loss of its activity is therefore best understood as failure to sustain the waking state, producing episodic loss

of wakefulness, rather than as a simple excess of a sleep-generating signal. That distinction matters because it prevents a common conceptual reversal: a wake-stabilising system can be clinically decisive even when it is not the system that first turns wakefulness on. The detailed orexin framing, hypersomnolence syndromes and their management are covered in the Eating and sleep-wake disorders notes.

The named anatomical landmarks provide the circuit vocabulary. **Wake-side nodes include the tuberomammillary nucleus, locus coeruleus, raphe nuclei, pedunculopontine and laterodorsal tegmental nuclei, lateral hypothalamus and ventral tegmental area; the ventrolateral preoptic area is the sleep-side node**. For the psychiatrist, their value is organisational: distributed brainstem, hypothalamic and forebrain-linked systems regulate a global behavioural state. The complete arousal-system anatomy, flip-flop teaching, sleep architecture and hypnogram should be revised from the Eating and sleep-wake disorders notes rather than reconstructed here.

RULE **Think in physiological systems before diagnostic labels.** Ask whether the dominant disturbance concerns state maintenance, ventilation, circadian alignment or the expression of state in measurable channels.

12.2 From circuits to observable physiology

A patient can appear asleep while the underlying physiology is fragmented, mis-timed or behaviourally dissociated. Conversely, a patient may report little or no sleep despite a recording that demonstrates an organised sleep pattern. The examiner therefore needs two levels of description. The clinical level contains the history of sleep opportunity, timing, continuity, daytime consequence, witnessed behaviour and breathing. The physiological level asks what the brain, eyes, muscles, cardiorespiratory system and body were doing at the same time. Neither level replaces the other.

This distinction explains why normal sleep architecture is foundational even when it is not re-taught in this chapter. Architecture supplies the reference pattern against which abnormal breathing, loss of expected muscle atonia, unusual movements, behaviours and state transitions become legible. It also prevents a category error: a symptom such as fatigue, poor concentration or nocturnal fear is not itself a physiological state. It may arise from disrupted sleep, circadian misalignment, impaired ventilation, a nocturnal event, a psychiatric disorder, medication effects or more than one process acting together. The sleep study can answer selected physiological questions, but it does not replace the formulation.

MNEMONIC

STATE - BREATH - CLOCK - TRACE - CONTEXT

Build the neurological account in that order: identify the state-control question, check whether ventilation disrupts sleep, decide whether timing is aligned, select the trace that can capture the event, and interpret it in clinical context.

The sequence is particularly useful in psychiatry because the presenting complaint often sits downstream from the disturbance. Irritability, cognitive inefficiency, mood symptoms or unusual nocturnal behaviour may bring the patient to a psychiatrist, while the decisive physiological abnormality lies in ventilation, timing or state-linked motor control. The correct response is neither to psychologise the complaint nor to treat the trace as the whole patient. It is to connect symptom, timing, witness account and physiological hypothesis.

12.3 Sleep-related breathing: one mechanism, several diagnoses

The shared neurological burden. **Obstructive sleep apnoea, central sleep apnoea and sleep-related hypoventilation share impaired ventilation during sleep, commonly accompanied by intermittent or sustained hypoxaemia and sleep disruption.** Their immediate mechanisms are not interchangeable, but the shared consequence is crucial for psychiatric reasoning. Repeated disruption of sleep continuity and impaired oxygenation can alter daytime alertness and cognitive or emotional functioning. A psychiatric presentation may therefore be downstream from a physiological sleep disorder even when the patient does not volunteer a sleep complaint.

Breathing history should remain mechanistic. Ask whether the account suggests obstruction, loss of respiratory drive, inadequate ventilation or a mixed picture, and whether the suspected disturbance is episodic or sustained. Witnessed nocturnal events, daytime consequences, medical context and medication exposure help form the pre-test hypothesis. The formal diagnostic indices, screening instrument, severity bands, positive-airway-pressure treatment and psychotropic-choice rules belong in the Eating and sleep-wake disorders notes and should not be inferred from symptoms alone.

The association with psychiatric illness is clinically important but causally untidy. In a sleep-apnoea cohort, **mood disorders, anxiety and posttraumatic stress disorder, psychosis and dementia were significantly more frequent than in the non-sleep-apnoea comparison group.** The direction remains contested. Sleep disruption and hypoxaemia may contribute to psychiatric symptoms, while psychiatric illness or its associated circumstances may predispose to apnoea. Comorbidity therefore supports active case-finding; it does not prove that either disorder caused the other.

> **PITFALL**
> Do not convert improvement-resistant mood or cognitive symptoms directly into a sleep-apnoea diagnosis. The correct move is to recognise the possibility, obtain the relevant sleep and breathing history, examine the medical context, and test the physiological hypothesis when indicated.

For an examination answer, keep the layers separate. State the **mechanism** as impaired ventilation with disturbed sleep and impaired oxygenation. Describe the **clinical expression** as nocturnal breathing disturbance with daytime neurobehavioural consequences. Then identify **confirmation** as a recording chosen to demonstrate the suspected respiratory disturbance during sleep.

12.4 Circadian timing: sleep ability versus sleep placement

Sleep generation and sleep timing are related but different questions. A person may possess the physiological capacity to sleep yet attempt it at a phase that is misaligned with internal timing or the external schedule. The **suprachiasmatic nucleus of the hypothalamus acts as the internal pacemaker for the sleep-wake cycle and also coordinates circadian hormonal and metabolic rhythms**. This gives circadian disturbance its broad clinical reach: the issue is not merely clock time on the bedside table, but alignment among internal physiology, environmental cues and required behaviour.

The need for entrainment becomes intuitive when external time cues are removed. Under such conditions, the internal cycle tends to lengthen to **24.5 to 25 hours**, rather than remaining locked to the environmental day. The discrepancy may seem small within one cycle, but without repeated resetting the internal phase drifts away from social time. External cues are therefore active synchronisers, not decorative accompaniments to an otherwise perfectly fixed clock.

Light is the central cue to understand. **Darkness promotes pineal melatonin synthesis and release so concentrations rise at night, whereas natural and artificial light suppress synthesis and release so concentrations fall during daylight**. The clinically important inference is that exposure and timing matter together. Light is not a generic sedative or stimulant; it is information supplied to a timing system. Exact phase-response rules, melatonin markers, circadian disorder subtypes, timed-light treatment and shift-work management are taught in the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes and the Eating and sleep-wake disorders notes.

TRAP **Do not describe every complaint of sleeping late as inability to sleep.** First distinguish failure of sleep generation from sleep occurring at an internally preferred but externally inconvenient phase.

12.5 Polysomnography as a set of physiological questions

Polysomnography is best learnt as simultaneous measurement, not as a mysterious machine that prints a diagnosis. A **full polysomnogram records cerebral activity through EEG channels, right and left ocular movements, chin and other muscle activity through EMG, oxygen saturation and other vital signs, together with video of movement, vocalisation and behaviour**. Each channel answers a different question, while simultaneity allows the reader to decide whether changes belong to the same event.

PHYSIOLOGICAL QUESTION	RECORDING DOMAIN	WHAT THE PSYCHIATRIST GAINS
What state is the cerebrum expressing?	EEG	A cerebral correlate against which other signals can be timed
Are characteristic eye movements present?	Right and left ocular channels	Ocular activity linked to the simultaneous cerebral and muscle pattern
Is muscle tone present, reduced or abnormally retained?	Chin, limb or other EMG	A measurable motor correlate rather than a witness description alone
Is oxygenation or other vital physiology disturbed during sleep?	Oxygen saturation and other vital-sign recordings	Temporal linkage between impaired oxygenation, physiological consequence and sleep disruption
What did the patient visibly do?	Time-linked video	Correlation of movement, vocalisation or behaviour with the physiological trace

The table should be read horizontally and vertically. Horizontally, each row moves from a clinical question to a signal and then to its interpretive use. Vertically, the channels create converging evidence about one period of recorded behaviour. A movement on video means something different when accompanied by expected muscle physiology than when it appears during a state in which that physiology should be absent. Likewise, an oxygen change is more informative when its relation to breathing and sleep disruption is visible.

IN INDIA

Access to comprehensive polysomnography may be uneven and cost can determine what is feasible. Referral should therefore state the physiological question clearly and include the relevant witness account, medication context and differential diagnosis. A costly recording without a defined question is less useful than a targeted study whose channels can adjudicate competing explanations.

12.6 Channel concordance: how the trace becomes neurophysiology

The diagnostic value lies in relationships among channels. **EEG separates sleep states, ocular channels show the eye movements, and EMG carries the muscle-atonia signal that distinguishes REM from NREM**. This makes a physiological absence, such as failure of expected REM atonia, recordable rather than merely descriptive. Video then shows the behaviour occurring at that moment. The study can therefore link a witnessed movement to cerebral state and muscle physiology instead of asking the clinician to choose between history and trace.

Interpretation requires concordance. A technically plausible signal should fit the neighbouring channels, the recorded video and the clinical question. Apparent movement may be irrelevant if it lacks the expected temporal relationship; an apparently normal channel may be unhelpful if the event was not captured; and an abnormal trace may be incidental to the presenting complaint. The report is strongest when it states what event occurred, which channels changed

together, which physiological state was present and whether that combination answers the referral question.

This is also why montage detail and formal staging cannot be replaced by a short list of electrodes in a psychiatry answer. Instrumentation determines what can be observed, but interpretation depends on validated recording practice and specialist scoring. Those details, along with sleep-stage scoring and the formal indices derived from the study, belong in the Eating and sleep-wake disorders notes. The raw electrophysiological principle of EEG is covered in the Epilepsy and psychiatry notes. Here, the transferable skill is to match each clinical uncertainty to the channel that could resolve it.

GOOD TO REMEMBER

When referring for a sleep study, write a testable question: suspected respiratory events, abnormal behaviour linked to REM physiology, repetitive limb movement, parasomnia or a possible sleep-only seizure. “Poor sleep” alone does not tell the laboratory what physiological event must be captured.

12.7 What polysomnography can confirm

The indication follows from the mechanism under suspicion. **Polysomnography can confirm sleep apnoea, REM sleep behaviour disorder, periodic limb movements and parasomnias; it is preliminary in narcolepsy assessment, for which a multiple sleep latency test can become diagnostic in the appropriate clinical setting**. This is not a list of interchangeable referrals. Each condition is supported by a different relation among cerebral state, muscle activity, ocular movement, respiratory physiology or recorded behaviour.

CLINICAL SUSPICION	PHYSIOLOGICAL TARGET	ROLE OF THE STUDY
Sleep-related breathing disorder	Ventilatory disturbance, oxygenation and associated sleep disruption	Confirms that the respiratory event occurs during recorded sleep and shows its physiological consequence
REM sleep behaviour disorder	REM state, muscle atonia and time-linked behaviour	Converts reported enactment into a state-linked motor finding
Periodic limb movements	Repetitive limb activity in relation to sleep	Measures the movement rather than relying on bed-partner description alone
Parasomnia	Behaviour in relation to physiological state	Helps classify the event and examine competing explanations
Suspected narcolepsy	Overnight context before daytime sleep-propensity testing	Provides the preliminary study; the subsequent test is interpreted with the clinical syndrome

The practical principle is to seek an event that the available channels can capture. If the question is breathing, the study must connect ventilation to oxygenation and sleep disruption. If the question is dream enactment, the study must connect behaviour to REM state and muscle tone. If

the question is a parasomnia or repetitive movement, video and motor channels must be interpreted against cerebral state. If the question is narcolepsy, overnight recording prepares the ground for the appropriate daytime investigation rather than replacing it.

Formal diagnostic thresholds, respiratory indices, movement indices, daytime test thresholds and actigraphy protocols are intentionally not reproduced here. They are covered in the Eating and sleep-wake disorders notes. In an answer focused on neurological foundations, marks come from stating what each test measures, why the measurement fits the hypothesis and how it changes the differential diagnosis.

12.8 Limited and negative uses: when a normal trace helps

Polysomnography has a limited role in routine insomnia evaluation, but a demonstrated normal sleep pattern may assist indirectly with nocturnal panic attacks, paradoxical insomnia, factitious sleep-wake presentations and seizures occurring exclusively during sleep; it may also reveal an underlying disorder in sleep-maintenance insomnia. This is the psychiatrist's distinctive use of the investigation. A normal or organised trace is not “nothing”; it can weaken a proposed physiological explanation and redirect attention to mismatch between experienced sleep, observed behaviour and recorded state.

The word indirectly matters. A normal sleep pattern does not by itself diagnose panic, intentional symptom production or misperception of sleep. It changes the evidential balance when integrated with history and observed behaviour. Similarly, a recording that captures a stereotyped nocturnal event with an epileptic correlate can reclassify an apparent sleep disorder, but seizure diagnosis and EEG interpretation require the framework in the Epilepsy and psychiatry notes. A negative recording is also bounded by what was sampled: failure to capture the habitual event is not equivalent to proving that the event never occurs.

Before requesting a study for insomnia, ask:

- Is there a specific alternative disorder that the recording could demonstrate?
- Is the complaint primarily about sleep timing, sleep opportunity, perceived sleep or an observable nocturnal event?
- Would a normal trace meaningfully narrow the differential diagnosis?

If none of these questions identifies a measurable target, more instrumentation may add data without adding understanding. Careful history remains the gatekeeper.

12.9 Integrated clinical and examination method

The section can be consolidated into a single reasoning sequence. Begin with the complaint but translate it into physiology. Decide whether wake maintenance, sleep generation, ventilation, circadian placement or an abnormal state-linked behaviour is the leading problem. Use collateral history when the event occurs outside the patient's awareness. Then select the signal capable of testing the hypothesis and state what finding would support it. Finally, return the result to the clinical context rather than allowing the laboratory label to replace formulation.

A compact answer can follow this order:

1. **Name the system and disturbance:** connect arousal and sleep-side nodes, wake-maintaining orexin, the circadian pacemaker or ventilatory control to unstable wakefulness, disrupted sleep, impaired oxygenation, circadian misalignment or abnormal motor expression.
2. **Choose the measurement:** cerebral, ocular, muscle, respiratory, vital-sign or video channel, interpreted simultaneously.
3. **Define the boundary:** explain what the study can confirm, what it only supports and what a normal recording can or cannot exclude.

This approach also keeps cross-references disciplined. Architecture, formal sleep-disorder classification, diagnostic indices, circadian subtypes and treatment belong in the Eating and sleep-wake disorders notes. Molecular clock physiology and detailed melatonin markers belong in the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes. Seizure electrophysiology belongs in the Epilepsy and psychiatry notes. The neurological foundation retained here is the bridge from system to signal.

> Memorise the chain: state-control circuit, breathing and circadian timing generate the clinical hypothesis; simultaneous physiological channels make that hypothesis testable.

13 · Insomnia disorder: assessment, hyperarousal model and pharmacological and CBT-I management

Insomnia is not adequately understood as simple sleep loss. For the psychiatrist, its most useful neurological formulation is a disorder in which the capacity for sleep is repeatedly opposed by persistent arousal. This framing makes sense of a familiar clinical mismatch: the person reports exhausting nights, yet may present by day as tense, alert, mentally over-engaged, or worried rather than simply somnolent. The **daytime extension of hyperarousal** is therefore part of the syndrome's clinical logic, not merely a consequence of lying awake.

This section deliberately takes the neurobiological lens. Formal diagnostic assessment, insomnia severity instruments, behavioural treatment and the pharmacological treatment ladder belong in *the Eating and sleep-wake disorders notes*. They should be used to establish the disorder and plan treatment; the present account explains why the reported experience is biologically coherent. The model is neither a licence to dismiss distress as "anxiety" nor a claim that every person with insomnia has identical laboratory findings. It is a way of linking subjective wakefulness, physiological activation and the difficulty of disengaging from wake-promoting processes.

13.1 The clinical problem the model is trying to explain

Patients often describe a night in which tiredness and wakefulness coexist. They may want sleep intensely, monitor its absence, and experience the bed as a setting in which thought and bodily vigilance become more prominent. A purely deficit-based account would expect inadequate sleep to produce an uncomplicated increase in objective daytime sleepiness. The hyperarousal model instead proposes that the same processes that sustain wakefulness can remain active across the night and spill into waking life.

The relevant distinction is between fatigue and sleep propensity. Fatigue can be experienced as depletion, poor concentration, low motivation or a sense of having been worn down. Sleep propensity concerns the tendency to fall asleep under standardised laboratory conditions. In insomnia, the latter may be reduced, and patients **may not show increased objective daytime sleepiness** compared with people without sleep disorders. This does not invalidate the complaint. It asks the clinician to hear two true things at once: restorative sleep has been disrupted, and the arousal system may still be preventing the expected drift towards sleep.

RULE **Do not equate a complaint of poor sleep with a measured tendency to fall asleep.** The mismatch itself can be clinically informative when it is interpreted as possible persistent arousal rather than as proof that the patient is exaggerating.

This formulation is especially useful in psychiatric practice because the consultation can otherwise become falsely binary. Either the patient is said to have a “real” sleep problem, or the sleep problem is reduced to a secondary emotional symptom. Hyperarousal permits a more precise account: cognitive, autonomic and cortical activation can be present together, while their relative contribution differs across patients and across phases of illness. The clinician can then assess sleep without prematurely assigning a single cause.

GOOD TO REMEMBER

A patient who is exhausted but not obviously sleepy has not thereby failed to demonstrate insomnia. Ask how alertness, worry, bodily activation and effort to sleep change from evening into the next day; the pattern is often more informative than the adjective “tired”.

13.2 Cortical arousal in the sleeping brain

The most direct support for the model comes from sleep electroencephalography. Compared with good sleepers, people with insomnia have **greater high-frequency EEG power** both around sleep onset and during non-rapid-eye-movement sleep. The finding is consistent with increased cortical arousal persisting when the person is meant to be asleep. It matters because sleep is not defined only by stillness, closed eyes or elapsed time in bed. Its quality depends on the brain’s capacity to reduce wake-like processing.

More specifically, the electrophysiological signature described in insomnia includes globally increased fast beta activity in the **16 to 32 Hz** range and more centrally increased alpha activity in the **8 to 12 Hz** range during non-rapid-eye-movement sleep. The point for a psychiatrist is not to turn these values into a bedside test. They are research-level markers that explain why a person can appear to have slept yet experience sleep as light, unrefreshing or vulnerable to intrusion by thought and environmental signals.

Fast-frequency EEG components during sleep are interpreted as an electrophysiological indicator of arousal. This is a convergent observation rather than a diagnostic signature to be sought in every patient. The model does not require the clinician to infer a particular tracing from a history. Its value is conceptual: insomnia can involve wake-like cortical activity within a state that conventionally looks like sleep.

LEVEL OF OBSERVATION	WHAT THE HYPERAROUSAL FORMULATION CONTRIBUTES	CLINICAL READING
Subjective night	Wakefulness can feel effortful, vigilant and mentally active.	Explore the quality of being awake, not only the duration of wakefulness.
Sleep EEG	**Wake-like cortical activity** can persist at sleep onset and in non-rapid-eye-movement sleep.	Sleep may contain an arousal component.
Daytime state	**Daytime hyperarousal** may co-exist with a poor night.	Do not use apparent activation to dismiss impairment.
Laboratory physiology	**Bodily activation** can be described across autonomic, endocrine and metabolic domains.	Interpret findings as support for a systems model, not as a single confirmatory test.

TRAP **Do not present high-frequency EEG activity as though it were a clinical diagnostic criterion.** It is evidence for a model of arousal, while insomnia remains a clinical disorder requiring the assessment framework in the sleep-wake chapter.

13.3 Autonomic, endocrine and metabolic activation

The cortical account is strengthened by evidence beyond the EEG. Laboratory measures in insomnia have shown **generalised HPA-axis activation** and increased arousal, although the findings are not consistent across studies. That qualifier must remain in the clinician's account. The model is supported by a pattern of observations, not by a uniform biomarker profile that every patient must express.

Reported correlates include raised whole-body metabolic rate, altered heart-rate variability, and altered norepinephrine and cytokine secretion. Cortisol and adrenocorticotrophic hormone have been found to be raised particularly in the **evening and early sleep period**. The timing is clinically intelligible: a system that should be winding down may remain physiologically mobilised as the patient approaches sleep. It also cautions against an overly psychological account of arousal. Thoughts about sleep matter, but the model is not confined to thoughts.

These observations should not be used as a work-up checklist. In ordinary psychiatric practice, they do not license indiscriminate endocrine, autonomic or inflammatory testing for insomnia. Their significance is explanatory. They show that the clinical experience can involve distributed activation across brain, endocrine and bodily systems. This is why the patient may describe the night as both mentally and physically "switched on", even when there is no external demand requiring alertness.

Cortical	Autonomic	Endocrine	Metabolic
STATE REGULATION	BODILY STATE	PHYSIOLOGY	ENERGY USE

The word "hyperarousal" is therefore best read as a systems-level shorthand. It refers to an interaction between cortical activation, bodily activation and attentional engagement with sleep. It does not imply that all insomnia is caused by stress, that the patient is voluntarily maintaining

wakefulness, or that a normal laboratory measure excludes the condition. The measured domains offer mutually reinforcing perspectives, while their inconsistency reminds us to avoid turning a model into an overconfident explanation of an individual case.

There is a further clinical advantage to this systems account. It helps explain why interventions aimed at only one surface feature may feel incomplete to a patient. Someone may reduce time awake in bed but still describe sleep as effortful; another may feel less worried about sleep but remain physically keyed up at night. Such observations do not prove a discrete biological subtype. They simply remind the clinician that cortical, cognitive and bodily components can travel together without being identical. A formulation should therefore be revised by the patient's longitudinal report rather than defended as a fixed explanatory label.

For the same reason, physiological language needs to be used with restraint in the consultation. Explaining that insomnia can involve activation across several systems can reduce shame and make the complaint easier to discuss. Explaining it as a measured endocrine abnormality in a particular patient, when no such measurement exists, is false precision. The senior clinical move is to use biology to deepen empathy and sharpen questions, not to convert a useful model into a laboratory diagnosis.

13.4 Failure of deactivation rather than absence of sleep capacity

Neuroimaging has encouraged a particularly useful reframing. The proposed problem is an **inability to switch off arousal-related circuits** rather than an inability to switch on sleep-related circuits. This is framed as suggestive evidence rather than settled fact, but it is clinically generative. It shifts the question from "Why cannot this brain make sleep?" to "What is keeping wake-promoting activity from relinquishing control?"

That distinction avoids a common therapeutic misunderstanding. If the difficulty is conceived only as an absence of sleep, treatment can become a race to impose sleep from outside. If it is conceived as persistent activation, the psychiatric task becomes broader: reduce the conditions that reinforce alertness, address maladaptive attention to sleep, and use treatment plans that are proportionate to the clinical setting. The detailed behavioural and pharmacological approaches are deliberately considered in *the Eating and sleep-wake disorders notes*; this neurological formulation explains their target without reproducing their protocols.

The word "switch" should not be understood as a simple on-off button. It is a clinical metaphor for state regulation. Arousal ordinarily supports attention, threat detection, thought and action. In insomnia, those functions may be insufficiently downregulated when sleep requires disengagement. The resulting experience is often paradoxical: the person can feel unable to stop trying to sleep precisely because the effort itself maintains a state of monitoring and readiness.

MNEMONIC

OFF

Observe the mismatch; **F**ind what maintains wakefulness; **F**ormulate the pattern from night into the waking hours. "OFF" keeps the clinical focus on disengagement rather than on forcing sleep.

13.5 A whole-day formulation

The model becomes most clinically useful when it is extended beyond the night. Some patients with insomnia are **hyperaroused and anxious in the daytime** and, despite poor sleep, do not necessarily feel sleepy. This is not a universal presentation and should not be used as a required feature. It is an explanatory pattern that helps the psychiatrist understand why a patient may be distressed, irritable, cognitively overactive and physiologically tense while also reporting a severe sleep complaint.

In the interview, this supports a timeline rather than a single question about "hours slept". Ask about the transition into evening, the onset of monitoring or worry, the bodily sense of activation in bed, the quality of nocturnal wakefulness, and the patient's state after rising. The aim is not to convert the consultation into a laboratory experiment. It is to establish whether the story describes a state that fails to disengage across contexts.

This approach also improves differential clinical reasoning without trespassing into other sleep-disorder material. Marked involuntary sleep episodes, unusual nocturnal behaviours, breathing-related symptoms, circadian misalignment, medication effects and mood-state changes may each demand their own assessment path. Name the possibilities, but do not force them into the hyperarousal model. Their diagnostic and management frameworks are set out in *the Eating and sleep-wake disorders notes*, and mood syndromes are addressed in *the Mood disorders: classification, depressive and bipolar syndromes, neurobiology notes*.

> **PITFALL**
>
> Do not call all daytime activation "evidence of anxiety". A poor sleeper who remains alert may be showing the very paradox that the hyperarousal model predicts. Equally, do not use that model to bypass assessment for a mood episode, substance effect, or another sleep-wake disorder.

13.6 How the model should shape assessment language

The assessment contribution of this model is qualitative. It teaches the psychiatrist to distinguish sleep opportunity from sleep capacity, tiredness from sleepiness, and a subjective sense of poor sleep from the physiological processes that may sustain it. It also invites attention to the patient's relationship with wakefulness: is the person frightened by it, struggling against it, monitoring it, or interpreting it as evidence of impending failure? These are clinically actionable formulations, even though the treatment methods themselves are owned elsewhere.

Use tentative language. "The history may suggest persistent arousal" is better than "the patient has excessive cortisol" or "the EEG is hyperaroused." The former is a formulation; the latter makes an unmeasured biological claim. This distinction protects the patient from false precision and protects the clinician from treating a group-level research finding as though it were an individual laboratory result.

The model is also a guard against moralising. Hyperarousal is not laziness, weak will or a simple failure of sleep hygiene. It describes a state in which brain and body may continue to privilege alertness. That formulation validates the patient's suffering while preserving clinical curiosity

about psychiatric comorbidity, substances, physical illness and other sleep conditions. It is more disciplined than assuming a single cause and more humane than treating insomnia as a complaint that should resolve through exhortation.

At the end of a psychiatric interview, the hyperarousal formulation should help organise the record, not replace diagnostic work. Useful language identifies what has been reported, what remains to be clarified, and which alternative explanations need their own pathway. It prevents two opposite errors: giving a biological explanation before the sleep history has been taken, and withholding a biological explanation until an investigation proves it. In most cases the formulation remains provisional and is tested against the course of symptoms and response to appropriately chosen treatment.

- Document the patient's account of mental and bodily activation separately from the complaint of inadequate sleep.
- Record daytime fatigue and daytime sleepiness as related but non-interchangeable experiences.
- State competing psychiatric, medical, substance-related and sleep-wake explanations explicitly rather than allowing "hyperarousal" to become a catch-all conclusion.

This way of writing the formulation also improves communication with the patient. It acknowledges that the symptom is real, makes uncertainty visible, and avoids making a speculative neurobiological explanation sound like a verdict. It leaves room for the person's own account of stress, mood, routines, bodily discomfort and the meanings attached to a poor night. Those details are not distractions from neuroscience. They are the clinical contexts in which persistent arousal is recognised and addressed.

13.7 Limits of the model and the bridge to treatment

Hyperarousal is an organising model, not a complete final mechanism. It remains under investigation whether the principal driver is **cortical hyperactivity or excess wake-stabilising orexin**. That uncertainty is important. It prevents a false impression that a single neurotransmitter account has solved insomnia, and it creates a principled bridge to the neurobiology of wake-stabilising systems considered with hypersomnolence disorders.

A practical formulation therefore has several layers. At the experiential level, the patient feels unable to disengage from wakefulness. At the cortical level, fast-frequency activity during sleep supports the possibility of residual arousal. At the bodily level, autonomic, endocrine and metabolic observations suggest that activation may not be confined to conscious thought. At the systems level, imaging suggests deficient downregulation of arousal-related circuitry. None of these layers should be over-read in an individual patient, but together they make the syndrome more intelligible.

Management should follow a full sleep-wake assessment and the appropriate treatment framework, not an isolated biomarker or an appealing neurobiological story. The psychiatrist's neurological contribution is to recognise persistent arousal as a plausible target, to avoid

mistaking fatigue for sleepiness, and to communicate that insomnia can be biologically real even when it does not look like ordinary sleep deprivation. For diagnostic assessment, cognitive behavioural treatment for insomnia and pharmacological management, see *the Eating and sleep-wake disorders notes*.

> Insomnia is best understood as a possible failure to downregulate arousal across the night and, in some patients, into the day.

14 · Narcolepsy and hypersomnolence disorders: hypocretin, cataplexy and management

The clinical question is not simply whether a patient sleeps too much. Narcolepsy is a disorder of unstable boundaries between wakefulness and sleep, and the orexin system gives psychiatrists a neurological way to understand that instability. This perspective prevents an unhelpful split between "neurological" symptoms and "psychiatric" experiences. A person may describe frightening imagery, immobility, irresistible sleepiness, emotional consequences, or apparent changes in behaviour. The task is to establish when the experience occurs, what state transition surrounds it, and whether it belongs to a coherent sleep-wake syndrome.

The marks in this area lie in linking a state-boundary symptom to its physiology without claiming more diagnostic certainty than the history supports. Detailed diagnostic criteria, formal sleep testing, disease classification and the treatment ladder belong in *the Eating and sleep-wake disorders notes*. The neurological foundation here is the wake-stabilising system, the logic of rapid eye movement sleep intrusion, and the psychiatric interview consequences of taking sleep timing seriously.

14.1 The orexin system: one system, two names

Orexin and hypocretin are two names for the same peptide neurotransmitter system. The terminology reflects a historical naming collision: one line of work emphasised the system's hypothalamic location, whereas the other drew on appetite in the name. Thus hypocretin **1** and **2** correspond to orexin A and B. The practical point is simple: a guideline, laboratory discussion, or medicine name may use either vocabulary. They should not be read as rival transmitters or as separate diagnostic entities.

The system should be conceived as a stabiliser rather than as a generic switch that merely starts wakefulness. During waking, **tonic firing of orexin neurons** maintains arousal. That is a useful correction to a crude "awake versus asleep" model. Wakefulness is an active, sustained state that must hold attention, muscle tone, perception and behaviour together while competing sleep processes remain possible. When its stabilising influence is lost, the clinical picture can include porous transitions between states rather than a simple excess of ordinary tiredness.

This system also matters beyond sleep. Orexins are implicated in **feeding and reward**, alongside other behaviours. That broader reach explains both the appetite-derived term and the psychiatric interest in the system. It does not license a reductionist reading of eating, motivation or reward

symptoms as proof of an orexin disorder. Rather, it reminds the clinician that arousal is embedded in behavioural regulation and that sleep histories deserve a place in a psychiatric formulation.

A useful formulation therefore has two levels. At the biological level, orexin helps maintain a coherent waking state. At the clinical level, the person may experience the failure of that stability as sudden sleepiness, a mismatch between bodily state and intended activity, or events that appear at the boundary between sleep and waking. This translation prevents a common conversational mistake: asking only whether the person is tired. It is more informative to ask how alertness behaves across the day, what situations expose the difficulty, and whether sleep-related experiences intrude into the waking account.

Wake	Feeding	Reward	Language
OREXIN STABILISES AROUSAL	A RECOGNISED WIDER FUNCTION	A RECOGNISED WIDER FUNCTION	OREXIN AND HYPOCRETIN ARE INTERCHANGEABLE

RULE **Read orexin as the system that steadies waking.** This makes the symptoms of narcolepsy intelligible as instability at a state boundary, not as a failure of will or a purely subjective complaint.

14.2 Hypothalamic source and the logic of cerebrospinal fluid measurement

The relevant cells are **lateral hypothalamic neurons**. They produce hypocretin and release some of it into cerebrospinal fluid. This anatomical fact explains why **cerebrospinal fluid hypocretin measurement** is biologically meaningful: the assay is not an indirect psychological marker but samples a peptide that reaches that compartment.

For a psychiatrist, the important distinction is between the rationale for a test and the test's place in a diagnostic pathway. A result is interpreted with clinical phenotype and specialist sleep assessment; it does not replace a careful history. The numerical laboratory threshold, the associated genetic marker, and the formal testing framework are deliberately not reproduced here. They are taught in the diagnostic section of *the Eating and sleep-wake disorders notes*. Keeping this boundary clear avoids the common error of treating a biomarker as a detached screening answer.

The hypothalamic location is also conceptually useful. It joins sleep-wake regulation to a region familiar to psychiatrists through motivation, appetite and endocrine discussion. Yet localisation should not become a shortcut to diagnosis. The clinician's job remains syndromic: identify persistent sleepiness, reconstruct state-linked events, and decide whether the whole account warrants formal sleep evaluation. A test can refine a well-framed question; it cannot repair a history that has not distinguished sleepiness from fatigue, low drive, medication-related sedation, or a fluctuating medical state.

That distinction has direct consequences for the referral letter. A useful referral does not merely state that a patient is “sleepy”. It describes the duration and pattern in ordinary language, the timing of episodes, sleep-related imagery or paralysis if present, the medications and substances

that may alter alertness, and the functional situations in which reduced alertness matters. It also records the psychiatric context without allowing it to become a catch-all explanation. This is not administrative detail. It lets the receiving clinician see why a sleep-wake mechanism is being considered and what alternative explanations still need attention.

GOOD TO REMEMBER

A cerebrospinal fluid assay has a direct physiological rationale because hypocretin produced in the lateral hypothalamus is secreted into cerebrospinal fluid. The result belongs within specialist diagnostic reasoning, not in isolation.

14.3 Narcolepsy as a hypocretin-deficit model

The discovery of a **selective loss of hypocretin-producing neurons** changed the explanatory centre of narcolepsy with cataplexy. The syndrome is no longer best understood as an undifferentiated complaint of daytime sleepiness. The model connects a discrete arousal system with the instability of wakefulness and with sleep phenomena appearing in an inappropriate state. It also gives a coherent account of why symptoms that might initially seem unrelated can be elicited in the same interview.

The mechanism of neuronal loss **may be autoimmune**. The hedge matters. “May” identifies a plausible pathophysiological account without turning a developing inference into an established individual cause. In an examination answer, this is the appropriate level of certainty. In clinical work, it protects against the equally unhelpful habits of presenting an aetiological theory as a confirmed diagnosis or dismissing the biological model because the patient also has understandable psychological distress.

Hypersomnolence is especially easy to flatten into an everyday word. The interview should instead ask what the patient means by sleepiness. Is there a pull towards sleep, an inability to remain alert, a loss of energy, slowed initiation, or sedation after a substance or medicine? These descriptions can overlap in ordinary language, but they direct different lines of enquiry. A temporal map of sleep, waking, experiences at the edge of sleep, medication exposure and functional consequences often provides more neurological information than a broad question about “tiredness”.

Psychiatric comorbidity or emotional reactions do not invalidate this enquiry. Fragmented function, fear of symptoms, shame about apparent sleepiness, and interpersonal misunderstanding may all shape the presentation. The neurologically disciplined move is to preserve phenomenology before assigning explanation. If the report is state-bound and internally consistent, it deserves sleep-wake formulation even when the consultation began in a psychiatric clinic.

This approach is particularly valuable when the narrative sounds inconsistent at first hearing. A patient may be able to describe an episode in striking detail and nevertheless be unable to say why it occurred. That is not evidence of fabrication. State transitions are brief and subjective, and recall can be shaped by fear, embarrassment and the language available to the patient. The

interviewer's role is to reconstruct the sequence patiently: the preceding state, the event itself, the recovery, and the surrounding sleep pattern. A coherent temporal account carries more weight than an initially dramatic label.

PITFALL

Do not use a plausible psychological explanation to close the sleep history. A patient can have distress, altered mood or medication exposure and still require a careful account of whether experiences cluster around transitions between wakefulness and sleep.

14.4 REM intrusion: the unifying neurological frame

REM sleep phenomena intruding into wakefulness provide the unifying physiological frame for cataplexy, sleep paralysis and sleep-related hallucinations in narcolepsy. This is a mechanism, not a licence to diagnose narcolepsy from any one experience. It explains why a wakeful person may report a bodily state resembling REM-related atonia or imagery resembling dream experience, while still retaining enough awareness to describe the event.

Cataplexy is best held here as a state-boundary phenomenon within the syndrome. Its detailed phenomenology, diagnostic role, and treatment are covered in *the Eating and sleep-wake disorders notes*. The essential neurological insight is that it is not best conceptualised as ordinary weakness, a voluntary behaviour, or a global loss of consciousness. The same frame makes sleep paralysis intelligible as the motor component of REM physiology persisting briefly into waking.

This formulation gives psychiatrists a disciplined route through seemingly dramatic accounts. First ask whether the episode sits at a sleep boundary. Then ask what consciousness, movement, imagery and recall were like during it. Finally, seek the wider sleep-wake context rather than deciding from an isolated symptom. The sequence is useful because it resists two opposite errors: medicalising every unusual nocturnal experience and treating a state-linked biological event as evidence of primary psychosis.

MNEMONIC

State, Sequence, Syndrome

State: locate the event near sleep or waking. **Sequence**: separate movement, imagery and awareness. **Syndrome**: decide only after the wider sleep-wake account has been taken. This is a history structure, not a diagnostic criterion.

TRAP **Calling every paralysed or hallucinatory episode narcolepsy.** These events can be compelling, but their presence alone does not establish the syndrome.

14.5 State-boundary hallucinations and the psychiatric interview

Sleep-related hallucinations typically occur at sleep onset or on awakening. The labels **hypnagogic and hypnopompic** are worth using because they force the clinician to document timing rather than merely record “hallucinations”. Their characteristic quality can be vivid and image-like; the experience may be immobile, frightening, and last for minutes. It commonly

recedes when a light is turned on. Those details point to a state transition and are often more discriminating than the emotional intensity of the account.

The interview must neither minimise nor overinterpret the experience. A frightened patient may understandably describe it in the language of an external presence or threat. The clinician can validate the fear while asking precise questions: Were the eyes open? Was the event linked to falling asleep or waking? Could the person move or speak? Did orientation return with environmental change? Was there a comparable experience in the fully awake daytime state? The purpose is phenomenological clarification, not adversarial cross-examination.

Collateral information can be useful when episodes disrupt shared sleep, work or family routines, but it should supplement rather than replace the patient's account. It may clarify observed behaviour, recovery, or the practical effect of sleepiness. It cannot by itself determine the internal experience of imagery, paralysis or awareness. The same caution applies to labels supplied by family members. A description such as “fainting”, “fits”, “possession”, or “behavioural problem” is a starting point for inquiry, not a conclusion. The psychiatric interview remains a careful exercise in translating lay descriptions into state, sequence and context.

It is reasonable to infer that state-bound timing, vivid visual quality, preserved recall, immobility and resolution with light should widen the differential beyond a primary psychotic process. That inference is not a claim that narcolepsy is commonly mistaken for any particular psychiatric disorder, nor does it make the distinction automatic. Psychiatric formulation remains necessary, especially when symptoms are not consistently tied to the sleep-wake boundary or when other psychopathology requires its own assessment.

HISTORY AXIS	STATE-BOUNDARY SLEEP EXPERIENCE	PSYCHIATRIC ASSESSMENT IMPLICATION
Timing	Sleep onset or awakening	Document the transition, not only the content
Perceptual form	Vivid, image-like and immobile	Clarify form before assigning meaning
Course	Can persist for minutes	Ask for sequence and recall
Environment	May disappear with light	Ask what changed the episode
Emotion	Can be frightening	Validate distress without letting fear determine the diagnosis

14.6 Sleep paralysis is informative but not diagnostic

Sleep paralysis is a brief intrusion of normal REM atonia into wakefulness. This gives it a neurological explanation without making it pathognomonic. Sleep paralysis and hypnagogic or hypnopompic hallucinations are **nonspecific symptoms**: they may occur alone or alongside other sleep disorders. The central skill is therefore calibration. An isolated report warrants careful exploration, but it cannot carry the whole diagnosis.

The distinction between "compatible with" and "diagnostic of" is a high-yield habit across neurology and psychiatry. Patients understandably seek a single label for an alarming event; clinicians are tempted to provide one, particularly when a symptom resembles a familiar textbook syndrome. Better practice is to identify the experience accurately, assess its frequency and context without inventing a threshold, and then decide whether the broader account supports referral for specialist evaluation.

When questioning, make room for ordinary and culturally shaped language. "Could not move" may refer to true bodily immobility, fear, dissociation, exhaustion, or a metaphor for helplessness. "Saw something" may denote imagery at a state boundary, an illusion, a dream remembered on waking, or an experience requiring a different psychiatric formulation. Clarification protects against both invalidation and premature diagnostic closure. It also helps the patient feel that the clinician is taking an apparently strange experience seriously without assuming its cause.

- Anchor the account to the transition into sleep or return to waking.
- Separate immobility, imagery, awareness and memory of the event.
- Ask whether the event occurs in isolation or within a broader sleep-wake complaint.
- Record functional and safety consequences in the patient's own routine.

14.7 Management: psychiatric contribution within a sleep-specialist pathway

LOCUS: management usually begins in outpatient collaboration, where psychiatric and sleep histories can be integrated without forcing the patient to choose between services. Urgent assessment is appropriate when sleepiness or its consequences create immediate risk, but the durable work is longitudinal: establish a reliable description of the sleep-wake problem, review concurrent mental symptoms and medicines, and ensure that the patient reaches the diagnostic pathway able to test the working formulation. Formal sleep investigations, diagnostic thresholds and syndrome-specific treatment belong in *the Eating and sleep-wake disorders notes*.

FOCUS: in the near term, the focus is accurate symptom mapping and reduction of preventable harm from unrecognised sleepiness or destabilising symptoms. Over time, focus expands to the practical burden: stigma, work or study disruption, anxiety about episodes, family misunderstanding and the effects of a chronic explanatory label. Psychiatric review can distinguish secondary demoralisation from a separate mood, anxiety, substance-related or psychotic syndrome, while avoiding the false choice that only one domain can be "real".

MODUS: the psychiatrist contributes explanatory psychoeducation, a structured temporal history, medication review, support for adherence to specialist assessment, and treatment of independently established psychiatric conditions. Drug selection for narcolepsy and pharmacological suppression of REM-related symptoms require the disorder-specific guidance and specialist context given in *the Eating and sleep-wake disorders notes*. Psychological work is not a substitute for the neurological assessment, but it can help a patient recover agency, communicate symptoms accurately, and manage the social consequences of an episodic disorder.

Shared formulation is often the most useful intervention at the psychiatric interface. Write down which experiences are clearly state-bound, which are less clear, and which mental symptoms need parallel assessment. This gives the patient a way to report change without having to defend a single explanation. It also makes later review more coherent: a newly recognised sleep disorder may clarify part of the presentation, while a coexisting psychiatric condition may still need independent treatment. Neither discovery cancels the other.

Explain the model in plain clinical language: the problem can involve the brain's regulation of wakefulness and the boundary between waking and REM sleep. Such an explanation often reduces moralised interpretations such as laziness, weak motivation or deliberate behaviour. It should remain provisional until the diagnostic process is complete. The goal is not to promise certainty early; it is to offer a credible account of why careful assessment is worthwhile.

> The must-memorise principle: state-boundary symptoms are clues to a sleep-wake formulation, not stand-alone proof of narcolepsy.

Sleep, movement and neurodegeneration

Why it matters clinically. Abnormal movement around sleep is not merely a sleep complaint. It may expose failure of a brainstem motor-control system, signal an adverse effect or withdrawal state created by prescribing, or provide a window into neurological disease. The psychiatrist therefore has to identify the state in which movement occurs, reconstruct the medication timeline, and decide how much neurological meaning the phenomenon can legitimately carry.

This section follows that reasoning from rapid eye movement sleep behaviour disorder to restless legs syndrome. It does not repeat formal sleep-disorder criteria, polysomnographic indices, or disorder-specific treatment algorithms, which belong in the Eating and sleep-wake disorders notes. Instead, it develops the neurological lens: how normal REM immobility is generated and lost, why a secondary drug phenomenon must not be given the prognostic meaning of an apparently spontaneous disorder, how iron and dopamine may converge in restless legs syndrome, and why timing is more discriminating than the mere presence of an urge to move.

15.1 The REM atonia circuit

Normal REM sleep combines an activated dreaming brain with profound suppression of skeletal muscle activity. **REM atonia** is generated by perilocus ceruleus or subcoeruleus systems that produce functional quadriparesis, atonia and areflexia in limb and trunk muscles. This state-dependent immobility is protective because it prevents dream content from being translated into purposeful movement. The essential idea is not that motor programmes disappear during REM sleep. They are actively prevented from reaching the final common motor pathway.

The circuit has **two** major inhibitory routes. A glutamatergic projection from the **subcoeruleus area** reaches the medullary reticular formation, which then supplies GABAergic and glycinergic inhibition to motoneurons. A parallel glutamatergic projection recruits GABAergic interneurons that also inhibit motoneurons. Reduced noradrenergic and serotonergic drive contributes

additional disfacilitation. Thus, REM muscle silence reflects both active inhibition and withdrawal of excitatory support rather than a single descending switch.

RULE **Dream enactment is a failure of state-dependent motor inhibition, not evidence that dreaming has simply become more intense.** The neurological question is where REM motor gating failed; the psychiatric question is whether disease, medication or withdrawal altered that gating.

This systems view also explains why the phrase **REM sleep without atonia** is mechanistically informative. It names a dissociation between the REM state and the motor inhibition normally coupled to it. The concept is more precise than treating every nocturnal movement as equivalent, because different movements arise before sleep, during non-REM transitions, during REM sleep, or independently of sleep altogether.

15.2 What lesion evidence can and cannot localise

Experimental lesion work supports the circuit model. Small lesions in the pontine reticular formation lateral to the locus coeruleus, or in the medial medulla, can abolish REM atonia by removing inhibitory input to spinal motoneurons. This is powerful causal evidence that the motor block depends on a distributed pontomedullary pathway. It also warns against imagining a purely cortical explanation for complex behaviour during REM sleep.

Animal and human localisation do not carry equal certainty. Bilateral perilocus coeruleus lesions can produce a presumed REM sleep behaviour disorder phenotype in animals. In humans, however, there is only a suggestion of associations with diffuse hemispheric lesions, bilateral thalamic abnormalities or brainstem lesions. The clean experimental locus therefore becomes a broader and less certain clinical network. A patient with the phenotype cannot be reverse-localised to a tiny pontine lesion from behaviour alone.

The practical formulation should preserve levels of inference. Loss of atonia shows failure of a physiological function; it does not by itself prove a focal structural lesion. A compatible structural disorder may disrupt the network, a diffuse neurodegenerative process may impair it, and a medicine may alter state control without producing a visible lesion. Conflating those levels turns a useful circuit model into false localisation.

GOOD TO REMEMBER

When the history suggests dream enactment, document the sleep-state relationship and the exposure timeline separately. The same outward behaviour can reflect circuit disease, a medication effect or rebound after withdrawal, and those formulations do not carry the same neurological implication.

15.3 The prescribing paradox in REM sleep behaviour disorder

Psychiatrists meet a paradox: medicines that suppress REM sleep can nevertheless provoke the symptoms and signs of REM sleep behaviour disorder. One proposed explanation is that the medicine allows persistent muscle activity during the REM sleep that remains, thereby uncoupling REM generation from its motor inhibition. This mechanism is explicitly provisional,

but the clinical association is sufficiently important to make medication review part of neurological interpretation. **Drug-precipitated REM sleep behaviour disorder** is therefore a state-control problem produced in the context of exposure, not automatically evidence of an emerging primary neurodegenerative process.

Tricyclic antidepressants, monoamine oxidase inhibitors and selective serotonin reuptake inhibitors are implicated classes. The operational exposure list is wider: biperiden, caffeine, venlafaxine, selegiline and serotonin agonists may provoke or worsen the syndrome, alongside the implicated antidepressant classes. The list matters because the less obvious agents can be missed when the review is reduced to a generic question about antidepressants.

Withdrawal supplies the mirror-image mechanism. Episodes may appear during withdrawal from alcohol, meprobamate, pentazocine or nitrazepam. **REM rebound** after withdrawal may precipitate attacks, so an apparently contradictory history of symptoms emerging after a drug was stopped can still be pharmacologically coherent. Current exposure and recent cessation must be reconstructed as parts of the same REM-regulation history.

- Establish whether the behaviour was observed during sleep rather than inferred from a vivid dream report.
- Record current REM-suppressing or otherwise implicated substances, including agents outside routine antidepressant review.
- Ask about recent withdrawal as carefully as current use, because rebound can reverse the apparent temporal logic.
- Keep a secondary exposure formulation separate from the later question of neurological disease association.

> **PITFALL**
>
> Do not label every medication-associated episode as idiopathic REM sleep behaviour disorder and then attach a neurodegenerative prognosis to it. First decide whether the same phenomenology is better formulated as secondary to a medicine or withdrawal.

15.4 Nosology and the neurodegeneration window

The distinction between phenomenology and diagnosis is crucial. ***DSM-5*** diagnostic framing excludes episodes attributable to medication, illicit substances or medical conditions. This makes it likely that drug-precipitated episodes would not be accepted as the primary disorder because they are secondary, although that conclusion should remain cautious. A patient can therefore show the same REM motor disinhibition while receiving a different diagnostic formulation and a different prognostic interpretation.

At the other end of the reasoning pathway, REM sleep behaviour disorder can be associated with neurological disease. Polysomnographic documentation is ordinarily required in the diagnostic framing, with an exception when a **synucleinopathy** is already established. That exception is worth naming because it exposes the close clinical bridge between REM atonia failure and

neurodegeneration, but the detailed diagnostic rule belongs in the Eating and sleep-wake disorders notes.

The reported neurological associations include Parkinson disease, dementia, progressive supranuclear palsy, Shy-Drager syndrome and narcolepsy. This list is a window, not a conversion calculator. The associations justify neurological attention but do not assign an individual conversion probability. Medication-precipitated or withdrawal-emergent behaviour, by contrast, requires its secondary cause to remain visible in the formulation. The diseases themselves belong in the Dementias and neurocognitive disorders notes and the Neuropsychiatry of systemic and neurological disease notes.

QUESTION	WHAT IT TESTS	INTERPRETIVE LIMIT
Was REM motor inhibition lost?	State-dependent circuit failure	Does not localise a focal lesion by itself
Was an implicated medicine present?	A secondary exposure formulation	Does not justify primary-disorder prognosis
Did episodes follow withdrawal?	A rebound mechanism	Stopping a substance does not exclude causation
Is neurological disease already established?	The neurodegeneration bridge	Does not replace disease-specific assessment

15.5 Restless legs syndrome as a waking sensory-motor state

Restless legs syndrome is best understood as a sensory-motor state at the boundary of sleep rather than as movement occurring in sleep. Its **four** linked elements are uncomfortable paresthesias before sleep, an irresistible drive to move the legs, relief when the legs are moved, and return of the sensations when movement stops. The sequence matters more than any single symptom. Paresthesia generates an urge, movement temporarily resolves it, and immobility allows it to recur.

The urge arises primarily during rest, inactivity or bedtime, and the movements are not the movements of actual sleep. Patients may rise and walk, move their feet while seated, or make bicycling movements in bed to obtain relief. These actions can look like pacing or motor restlessness, but their meaning lies in the context and sensory consequence. Movement is purposeful self-relief, not merely an observed excess of activity. The accompanying figure sets these observations against their mimics; they need not be turned into another checklist here.

MNEMONIC

REST

Rest brings it on; **E**asing follows movement; **S**ensory discomfort drives the movement; **T**he discomfort returns when movement stops. The mnemonic preserves the linked sequence rather than reducing the syndrome to restless movement.

This is also why polysomnography is not the instrument that establishes the clinical syndrome itself. The relevant urge and relief occur while the patient is awake. Sleep investigation may answer a different question about movements occurring after sleep begins, but it cannot substitute for a careful sensory and temporal history.

15.6 Distinguishing restless legs syndrome from akathisia

The trap exists because the subjective architecture overlaps. People with restless legs syndrome, tics and **akathisia** may experience psychological discomfort when resisting an urge and relief after complying. An urge-relief sequence therefore does not discriminate by itself. Likewise, walking or repetitive leg movement can be shared outward behaviour. Diagnosis depends on the temporal geography of the complaint: where, when and under what state conditions the urge appears.

The circadian and rest-linked pattern is the strongest discriminator; drug-related akathisia and hypnic jerks lack the striking circadian pattern and periodicity of the restless-legs complex. The accompanying figure brings the discriminating axes and the worked medication case together without treating any isolated feature as decisive. Its deeper lesson is that exposure changes prior probability, while phenomenology determines what the observed movement actually represents.

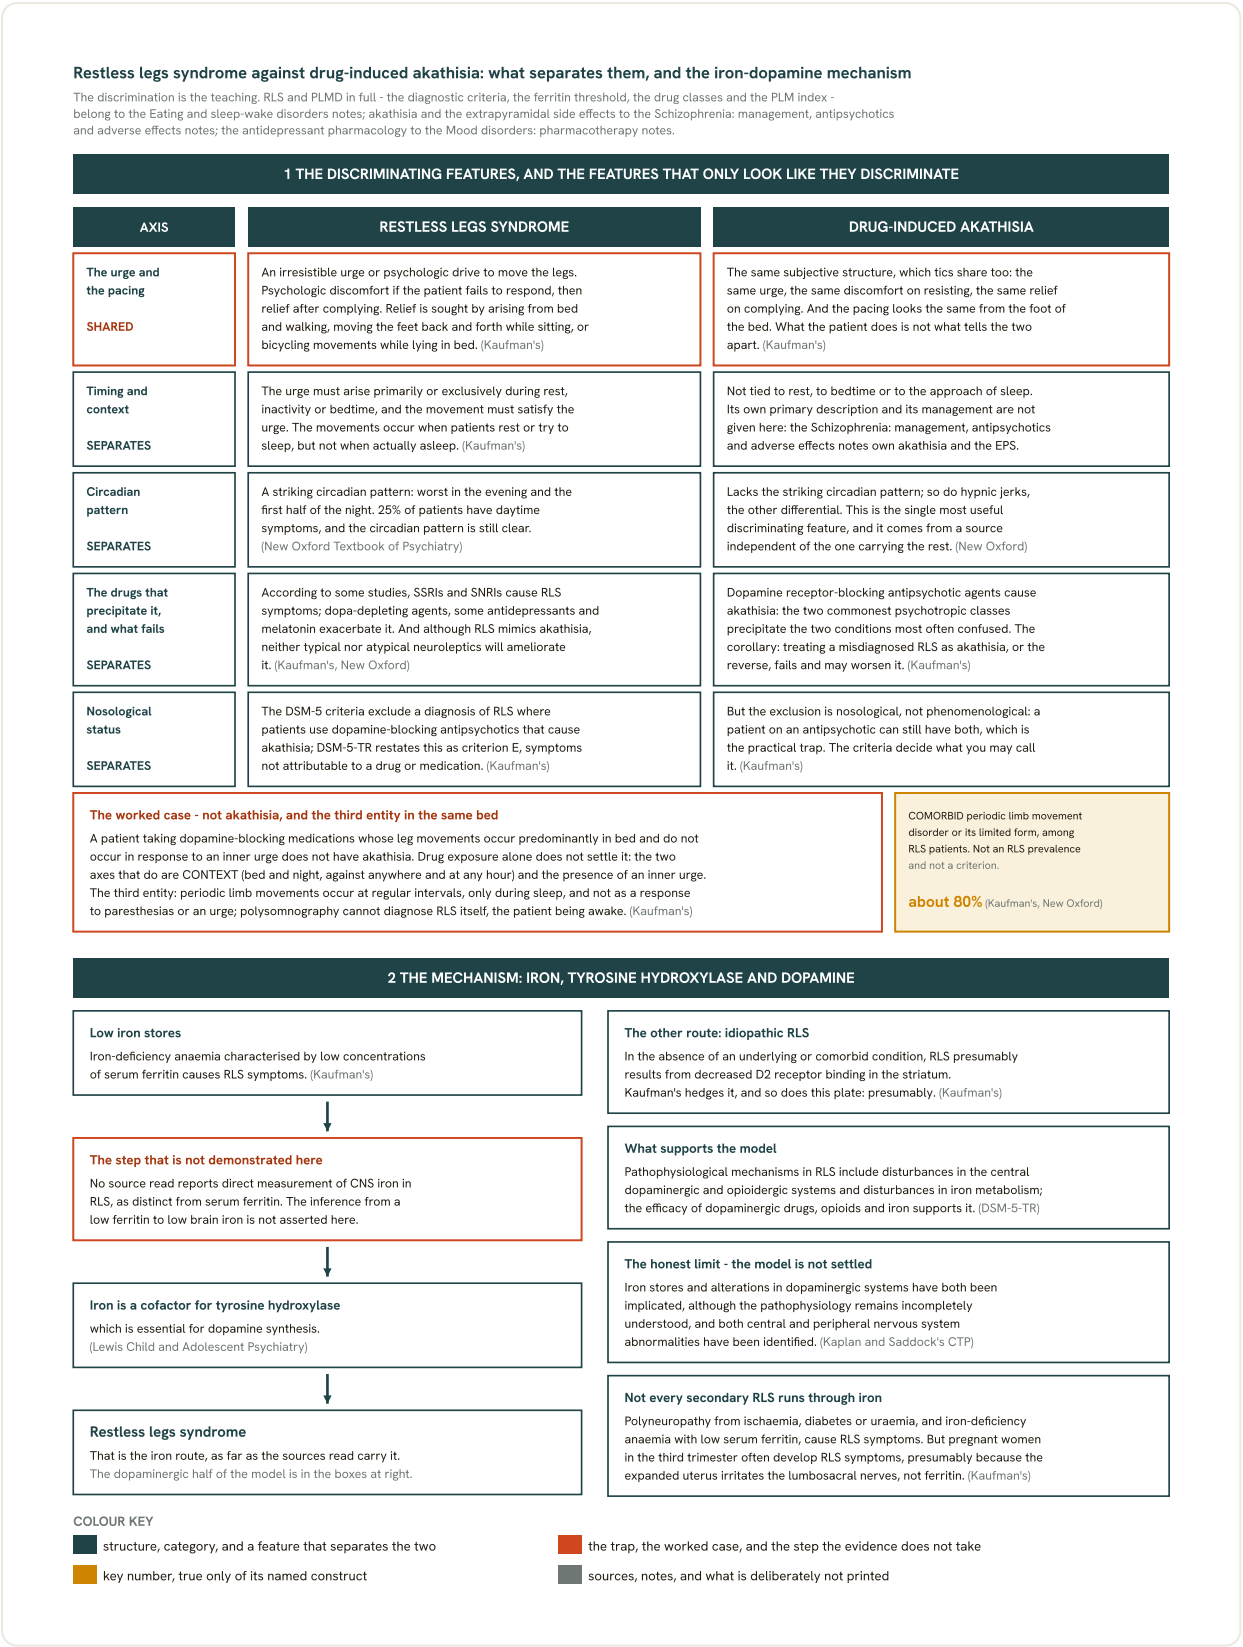


Fig 31.7 Restless legs syndrome against drug-induced akathisia. The urge and the pacing are shared, so neither discriminates; the context, the circadian pattern akathisia lacks, the precipitating drug class and the failure of neuroleptics to ameliorate RLS do. Akathisia appears only by contrast, cross-referred above. The iron chain is drawn with its hole visible: no source read measures CNS iron in RLS, so the ferritin-to-brain-iron step is named, not asserted, and the idiopathic D2 route is kept separate. Sources: Kaufman's Clinical Neurology for Psychiatrists, New Oxford Textbook of Psychiatry, Kaplan and Saddock's CTP, DSM-5-TR, Lewis Child and Adolescent Psychiatry.

TRAP **Do not diagnose akathisia from antipsychotic exposure plus leg movement alone.** Ask whether there is generalized inner restlessness, then map symptoms against rest, bedtime, circadian pattern and relief. Exposure is context, not the diagnosis.

Nosology creates an additional layer. DSM-5 excludes restless legs syndrome when symptoms are attributable to a dopamine-blocking antipsychotic causing akathisia. This is a rule about diagnostic attribution, not proof that a treated patient cannot have both syndromes. The clinician must first describe the phenomena accurately and then decide which diagnosis each phenomenon supports. The detailed examination and treatment of extrapyramidal syndromes belong in the Schizophrenia: management, antipsychotics and adverse effects notes.

The figure also prevents a separate collapse between restless legs syndrome and **periodic limb movements**. The useful clinical move is to keep the patient's waking sensory account distinct from a bed partner's observation of movement during sleep. Coexistence does not make the phenomena interchangeable, and one label should not be allowed to erase the other history.

15.7 Iron, dopamine and the limits of a unifying mechanism

The iron-dopamine account is useful because it connects biology, secondary causes and pharmacological observations, but it must remain a model rather than a settled single lesion. Iron is a cofactor for **tyrosine hydroxylase**, which is essential for dopamine synthesis. Reduced iron availability can therefore provide a biologically coherent route to impaired dopaminergic function. In apparently idiopathic disease, decreased striatal **D2** receptor binding has been proposed, although this remains presumptive.

Convergent treatment response adds plausibility without proving a complete causal chain. Evidence implicates disturbances in dopaminergic and opioidergic systems as well as iron metabolism, supported by the efficacy of interventions acting on those systems. Yet iron stores and dopaminergic alterations do not explain every case, and both central and peripheral nervous system abnormalities have been identified. The mechanism should therefore guide questions and investigations while preserving uncertainty.

EVIDENCE STREAM	SUPPORTED INFERENCE	WHAT MUST NOT BE OVERCLAIMED
Iron biology	Iron supports an enzyme essential to dopamine synthesis	Not every case is caused by peripheral iron deficiency
Striatal finding	Reduced dopamine-receptor binding is a proposed idiopathic mechanism	The hypothesis is not a universal lesion
Pharmacological convergence	Dopaminergic, opioidergic and iron-related responses support involvement of these systems	Response does not establish a complete causal pathway
System distribution	Central and peripheral abnormalities have both been identified	Pathophysiology remains incompletely understood

Clinically, the model prevents a false dichotomy between a brain disorder and a systemic disorder. Peripheral neuropathic or metabolic disease can generate the symptom complex, while central state-dependent and dopaminergic mechanisms shape its expression. The sensible formulation can therefore contain a syndrome description, a possible secondary contributor and an explicitly hedged mechanism at the same time.

15.8 Secondary conditions, psychotropics and the depression loop

Secondary associations broaden the neurological work-up. Polyneuropathy related to ischaemia, diabetes or uraemia, and iron-deficiency anaemia with low serum ferritin, can cause restless legs symptoms. Pregnancy is a useful warning against forcing every secondary case through iron: during the **third trimester**, symptoms have been attributed, presumptively, to irritation of adjacent lumbosacral nerves by the expanded uterus rather than to low ferritin. The principle is to search for the relevant biological context without assuming a universal mechanism.

The psychiatric vicious circle. Serotonergic antidepressants may induce or aggravate restless legs syndrome, yet the syndrome itself may predispose to depression, and effective treatment of the restless legs symptoms can reduce depressive symptoms. A clinician may therefore intensify antidepressant treatment for low mood that is partly sustained by an unrecognized sleep-related sensory-motor disorder, worsening the contributor that needs attention. The review must ask whether the mood syndrome, the leg symptoms and the prescription are interacting rather than assuming a simple sequence in which depression came first and every later complaint is merely an adverse effect. General psychotropic pharmacology belongs in the Mood disorders: pharmacotherapy notes.

Medical comorbidities include iron deficiency, end-stage renal disease, neuropathy, fibromyalgia, diabetes, Parkinson disease and multiple sclerosis. These are associations, not proof of causation. In particular, this neurodegenerative comorbidity must not be rewritten as a prodromal relationship. That inference would require separate longitudinal evidence and would confuse coexistence with prediction.

15.9 Epidemiology and an integrated ward formulation

At least 5%	At least 50%
POPULATION PREVALENCE	FAMILY HISTORY

Restless legs syndrome occurs across ages, and familial cases tend to present younger. The high frequency of family history supports asking directly about similar night-time sensory-motor symptoms in relatives rather than relying on a volunteered diagnostic label. The accompanying figure explains why a bed partner's report and the patient's own account may seem to describe different phenomena: each observer has access to a different state and a different part of the history.

An integrated ward formulation should separate observation from inference. “Legs moving” is an observation, not a diagnosis. The history must then supply sensory quality, state, context, subjective drive and the effect of movement. The clinician's job is to resist collapsing those layers

into whichever diagnosis is most strongly suggested by the medication chart. This approach also permits coexistence: a medication-related movement syndrome does not make every later nocturnal movement part of that syndrome.

The same discipline applies to REM phenomena. Describe what was observed, establish its relationship to REM sleep through the appropriate sleep-disorder pathway, identify current or withdrawn precipitants, and only then consider the neurological disease window. RBD and RLS both reward state-based reasoning: movement has different meaning during REM sleep, before sleep while resting, and during sleep without subjective urge.

> Memorise this: in REM dream enactment, separate circuit failure from exposure and withdrawal; in leg restlessness, timing and state outrank the medication label.

Consolidate and apply

16 · Worked vignettes

Why vignettes matter clinically. Neurological assessment in psychiatry is not a ritual performed after the mental-state examination. It is a method for deciding whether an apparently psychiatric presentation can safely remain within a psychiatric formulation, whether it needs parallel medical enquiry, and how urgently the setting of care must change. The useful question is not, "What neurological disease does this patient have?" It is, "What feature of this presentation cannot be adequately explained until I examine, localise, investigate, or escalate?"

These vignettes rehearse the neurological lens rather than reproduce component teaching. Examination, localisation, imaging, electrophysiology, cerebrospinal-fluid work-up, headache syndromes, disorders of consciousness, and sleep disorders each have their own accounts elsewhere in this chapter. Read a vignette as a rehearsal of **problem representation**: compress the presentation into a meaningful sentence, identify the discordant feature, and choose the next defensible action. The diagnosis may remain open. That is often safer than premature certainty.

PRESENTATION	EXAMINATION	TRAJECTORY	CONTEXT
WHAT THE PATIENT REPORTS	WHAT THE CLINICIAN CAN VERIFY	WHAT HAS CHANGED OVER TIME	WHAT MAKES THE PATTERN PLAUSIBLE

16.1 New behavioural change with a discordant neurological finding

A patient is brought with irritability, disinhibition, reduced self-care, and a story from relatives that "the personality has changed." The interview can easily become organised around diagnosis, risk, and medication. It should also become organised around onset, tempo, functional decline, and the reliability of the account. Ask what was present before the behavioural change, what changed first, whether the decline was continuous or fluctuating, and whether the patient can give a coherent account of the change. The collateral history is not merely corroborative; it may be the only route to detecting loss of insight or a deterioration that the patient cannot describe.

The neurological lens begins when the bedside examination yields a finding that does not sit comfortably inside a purely affective, psychotic, or personality-based explanation. A deficit in a higher function, an asymmetry, a change in gait, an unusual movement, or altered arousal changes the task. Do not convert the finding into a named syndrome by pattern-recognition alone. Instead, construct a **sign-symptom mismatch**: the psychiatric symptom is real, but the objective sign demands a parallel neurological account. The point of this formulation is not to make every behavioural disturbance organic. It is to avoid allowing a familiar psychiatric label to silence a discordant observation.

The immediate response has distinct parts.

- Repeat and document the finding in plain examination language.
- Establish whether the finding is new, persistent, progressive, fluctuating, or observed only under one condition.
- Seek collateral about baseline function, safety, medication exposure, substance exposure, recent illness, and observed neurological change.
- Decide whether the setting permits observation and planned work-up, or whether the pattern requires urgent medical assessment.

Localisation is a reasoning discipline, not a parlour trick. The anatomical account and the formal localisation ladder belong in *the Functional Neuroanatomy and Neurophysiology notes*. The psychiatry-facing task is narrower: a reproducible focal or lateralised finding should make the clinician pause before attributing the entire presentation to a primary psychiatric disorder. If imaging becomes relevant, choose it through the examination-to-imaging bridge and consult *the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes* for modality and indication detail.

RULE **An objective neurological finding changes the formulation before it changes the diagnosis.** First establish what the finding is, whether it is reproducible, and what urgency it creates. Diagnostic naming can follow better evidence.

Notice the difference between respecting psychiatric phenomenology and being trapped by it. Disinhibition may be experienced by relatives as “mania”; impaired self-care may be read as “negative symptoms”; apathy may be labelled depression. Such descriptions remain clinically important, but they do not settle causation. The safest formulation holds both levels: the patient has a psychiatric presentation requiring care now, and the presentation contains a neurological feature requiring explanation.

16.2 A fluctuating confusional presentation on the psychiatric ward

A patient admitted for agitation becomes intermittently inattentive, disorientated, frightened, and difficult to engage. At lucid moments the patient appears more coherent, leading staff to debate whether the presentation is behavioural, psychotic, medication related, or deliberate. The mistake is to treat fluctuation as evidence against illness. Fluctuation is itself a finding. It directs

attention to arousal, attention, temporal course, exposure history, systemic illness, and the possibility that more than one process is occurring.

Begin with a **state examination**, not a long cognitive battery. Can the patient sustain attention long enough to understand a question? Does the level of contact change across the interview? Is the patient drowsy, over-aroused, internally preoccupied, frightened, or unable to maintain a coherent line of thought? Re-examination is essential because a single polished interview note can falsely imply stability. The nurse's contemporaneous observations, relatives' description of the baseline, and the sequence of medication administration may reveal the course more clearly than a single clinician encounter.

At this point, separate what belongs to descriptive psychopathology from what belongs to neurological assessment. The phenomenology of clouding, stupor, and related states is developed in *the Psychometry, Assessment and Descriptive Psychopathology notes*. Delirium and focal brain syndromes are developed in *the Stroke, TBI, delirium and focal brain syndromes notes*. This chapter's lens is to recognise that altered attention or arousal changes the reliability of every psychiatric inference made after it. A denial of suicidal intent, an apparently bizarre belief, or apparent non-cooperation cannot be interpreted in the usual way when contact itself is unstable.

The practical sequence is therefore one of containment and clarification. Protect the patient from preventable harm, obtain physiological assessment through the medical team, review recent changes and exposures, and use serial bedside observations to determine the direction of travel. If there is concern for seizure-related or encephalopathic activity, an electroencephalographic question may be appropriate. The technical foundations of that investigation belong in *the Epilepsy and psychiatry notes*; the psychiatric-work-up indications are covered in the chapter's electrophysiology section. A normal-looking interval at the bedside is not permission to discard the preceding pattern.

GOOD TO REMEMBER

When attention and arousal are unstable, the first clinical task is to establish the patient's current state and safety. A polished diagnostic formulation based on an unreliable encounter is less useful than a brief, accurate account of fluctuation and a clear escalation plan.

This vignette also tests professional communication. The handover should state what staff saw, what the patient could and could not do during assessment, what changed, and what needs repeating. It should not hide uncertainty under a label such as “behavioural.” That word explains nothing. It can become a form of diagnostic closure when the patient most needs medical and psychiatric teams to observe together.

16.3 Headache, visual complaint, and a psychiatric medication history

A patient in follow-up for a mood or psychotic disorder describes headache, nausea, visual disturbance, reduced concentration, and worry that medication is “causing brain damage.” The psychiatrist has two responsibilities at once. The first is not to dismiss symptoms because the patient has anxiety, affective illness, or a history of somatic preoccupation. The second is not to

turn every headache into an emergency neurological diagnosis. Good reasoning preserves both responsibilities by asking which features are new, what their trajectory is, whether the patient is functionally impaired, and what the focused examination contributes.

The symptom account should be translated into a **syndrome frame** before medication changes are made. Is headache the dominant complaint, or part of a broader alteration in vision, cognition, alertness, or motor function? Is the visual complaint vague discomfort, an intermittent perceptual symptom, or a sustained loss of function? Has the patient already altered medication without telling the team? Has a rescue medicine become part of the problem? These questions distinguish a medication discussion from an indiscriminate medication reaction.

Do not use the psychiatric diagnosis as a substitute for examination. A focused neurological assessment, including attention to visual function and relevant eye findings, may expose a reason to alter the urgency and pathway of care. The technical examination belongs in the chapter's neurological-examination section. The disorder-specific accounts of migraine, medication-overuse headache, and idiopathic intracranial hypertension belong in their respective sections. The medication choices used for migraine prophylaxis are not interchangeable with psychiatric prescribing: consult *the Mood disorders: pharmacotherapy notes* for general psychopharmacology and the chapter's headache sections for indication-specific decisions.

A useful interview structure is:

- Clarify the symptom narrative without accepting the patient's causal attribution uncritically.
- Map medication exposure, adherence, recent changes, and self-directed use of non-prescribed remedies.
- Examine for the neurological information that can change urgency.
- State explicitly whether the current plan is symptomatic support, planned investigation, or urgent referral.

TRAP **"It is probably anxiety" is not a neurological conclusion.** Anxiety may amplify symptom distress and coexist with a primary headache disorder or another medical process. It cannot replace a history, examination, and safety decision.

The same discipline protects against the opposite error: treating a distressing headache report as proof that the psychiatric medicine must be stopped immediately. Abrupt, poorly reasoned medication changes can worsen the psychiatric illness while failing to address the headache. Explain the uncertainty honestly, decide what must be assessed now, and coordinate with the relevant medical service. The patient hears something much more useful than reassurance: their symptom is being taken seriously without a false promise of instant attribution.

16.4 The sleep complaint that is not merely insomnia

A patient with depression reports poor sleep, daytime fatigue, irritability, impaired concentration, and reduced motivation. The tempting formulation is circular: depression causes poor sleep, poor sleep worsens depression, therefore treatment should proceed as usual. That formulation may be partly true, but it is incomplete until the sleep complaint is unpacked. "I do

not sleep" can mean difficulty initiating sleep, disrupted sleep, insufficient opportunity, irregular timing, involuntary sleep episodes, disturbed breathing, abnormal movements, frightening nocturnal events, or non-restorative sleep. Each history points the assessment in a different direction.

The key term is **sleep-wake phenotype**: the characteristic pattern of sleep, wakefulness, timing, behaviour, and daytime consequence. Elicit it in the patient's ordinary language before translating it into a diagnostic category. Ask the patient and, where possible, the bed partner what happens across the whole sleep period, not only at bedtime. Clarify work schedule, opportunity for sleep, stimulant and sedative use, changes after medication, and safety-sensitive daytime consequences. A patient who says "insomnia" may be using a familiar word for a problem that has a different clinical structure.

This is a place where psychiatrists should resist both overconfidence and unnecessary testing. The sleep-disorder assessment, behavioural treatment, medication choices, and formal sleep investigations are set out in *the Eating and sleep-wake disorders notes*. The neurology chapters of sleep explain the relevant physiology and investigation lens. In routine psychiatric care, the contribution is to recognise when the account has exceeded the boundaries of a straightforward insomnia formulation and to make a clear referral question rather than a vague request to "rule out sleep disorder."

The referral question should specify the observed problem and its consequence: unexplained sleepiness, suspected breathing-related sleep disturbance, abnormal nocturnal behaviour, prominent limb symptoms, unusual dream enactment, or concern about circadian misalignment. This protects the patient from having a sleep service receive only the psychiatric label. It also protects the psychiatrist from changing sedating or activating medicines before understanding whether the change is likely to clarify or further obscure the presentation.

> **PITFALL**
>
> Do not use the word "insomnia" as shorthand for every complaint of poor sleep. A complaint is the opening of the history, not its conclusion. Treating the label instead of the phenotype can conceal a condition whose daytime risk is the most important part of the story.

The psychiatric formulation can then become more precise. Sleep may be a maintaining factor, an adverse-effect concern, a marker of relapse, a co-occurring disorder, or a signal of another neurological process. These are different causal positions and lead to different questions. The value of the sleep lens is not that it produces a fashionable additional diagnosis. It prevents an undifferentiated symptom from governing treatment.

16.5 Apparent non-organic symptoms and the obligation to examine

A patient presents with weakness, altered sensation, a tremor, gait difficulty, episodes of unresponsiveness, or a dramatic movement that appears inconsistent during conversation. The consultation may contain obvious psychosocial stress, dissociation, prior psychiatric treatment, and a strong request for certainty. None of those features removes the obligation to conduct a

respectful neurological assessment. The correct aim is not to catch the patient out. It is to determine whether the observed pattern is internally coherent, whether it maps to a plausible neurological process, and whether there are positive clinical grounds for the formulation being considered.

Use **positive inconsistency** carefully. It means that the examination identifies a reproducible feature whose behaviour under appropriate testing is not compatible with the expected pattern of a structural deficit. It does not mean that a symptom is unusual, emotionally charged, fluctuating, or difficult to explain. Nor does it mean that normal investigations prove a functional presentation. A normal test can be valuable, but it answers only the question it was capable of answering.

The detailed neurological examination is owned by the chapter's bedside-examination section. The clinical account of dissociation and conversion belongs in *the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes*. This vignette teaches the interface: psychiatric formulation gains credibility when it is based on positive examination evidence, coherent history, transparent explanation, and continued readiness to revise. It loses credibility when it is merely the residual category after the clinician becomes impatient.

RULE **Functional symptoms require positive clinical reasoning, not a diagnosis of exclusion made in frustration.** The examination should be collaborative, adequately documented, and repeated if the clinical picture changes.

Communication is part of the neurological intervention. Explain what has been found without implying fabrication or triviality. Distinguish between "we have not yet found a dangerous explanation" and "there is nothing wrong." The latter is usually false and clinically corrosive. A clear formulation can validate genuine symptoms, preserve trust, and leave room for appropriate neurological review if the course changes.

Medication review also belongs in the formulation, but not as an automatic answer. Some medicines can alter movement, alertness, sleep, or subjective cognition. The assessment must first identify the phenotype of the complaint, establish temporal relations, and examine for objective change. For adverse-effect syndromes and their pharmacology, see *the Schizophrenia: management, antipsychotics and adverse effects notes*. This prevents a common drift in practice: calling a movement "functional" because it is unfamiliar, or calling it a medication adverse effect because a medicine is present.

16.6 Building an integrated neuropsychiatric plan

The final vignette is not a disease presentation. It is the common real-world situation in which a patient has an established psychiatric disorder, a new physical or neurological complaint, an uncertain medication history, and several teams who each see only part of the picture. The psychiatrist's task is to build a **parallel formulation**. One strand addresses current psychiatric symptoms, risk, capacity, engagement, and continuity of treatment. The other addresses the neurological question, its urgency, and the information required to resolve it. Neither strand should cancel the other.

In the outpatient LOCUS, the plan needs a named owner for each action, a time-bound review arrangement expressed in ordinary clinical scheduling, and clear return precautions. In an inpatient LOCUS, serial observation and nursing evidence become central, while the team must guard against diagnostic diffusion in which everybody assumes somebody else has examined the patient. In an emergency LOCUS, stabilisation and escalation come before elegant phenomenology. The same formulation can travel across settings, but the available evidence and acceptable delay differ.

A concise integrated note records the following information.

DOMAIN	RECORD	CLINICAL USE
Psychiatric presentation	Syndrome and immediate risks without overclaiming cause	Preserves psychiatric care
Neurological concern	Symptoms and signs rather than a speculative label	Makes the question visible to collaborators
Examination	Observations that alter the pathway	Explains the urgency decision
Next step	Investigation or specialist question and what would change management	Prevents unfocused referral
Communication	Plan for patient, family, nursing staff, and collaborating services	Creates continuity across settings

This structure prevents the false choice between "organic" and "psychiatric." Psychiatric illness can coexist with neurological illness. Medication effects can coexist with relapse. A neurological complaint can be genuine even when anxiety increases attention to it. The formulation should therefore be revisable. New collateral, a repeated examination, changing arousal, or an investigation result can appropriately move the balance of probability. Revisability is a strength because it documents what the clinician knew at the time and what evidence would make them act differently.

MNEMONIC

PAUSE

Before closing a neurological question in psychiatric practice, **P**icture the trajectory, **A**ssess the state and examination, **U**ncover exposures and collateral, **S**tate the safety and setting decision, and **E**scalate with a focused question. PAUSE is not a diagnostic algorithm. It is a safeguard against premature closure.

The cross-references are part of the plan, not an admission of weakness. Use *the Functional Neuroanatomy and Neurophysiology notes* for anatomy, *the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes* for imaging and foundational psychopharmacology, *the Psychometry, Assessment and Descriptive Psychopathology notes* for formal assessment, and *the Eating and sleep-wake disorders notes* for sleep-disorder management.

The psychiatrist retains responsibility for integration: making sure the neurological question is visible, the psychiatric treatment remains safe, and the patient is not passed between services without a coherent explanation.

> When a psychiatric presentation contains a neurological discordance, pause, examine, describe the uncertainty, and make the next action answer a focused clinical question.

17 · Consolidation and quick review

The neurological lens is a discipline of attention, not a separate specialty consultation performed in miniature. In psychiatric practice, its purpose is to decide whether the presenting disturbance can safely be understood as primarily psychiatric, whether an accompanying neurological process may be shaping it, and what immediate action follows from that distinction. The task is not to convert every consultation into an exhaustive neurological examination. It is to recognise when the mental state, bodily signs, time course, and functional change do not form a coherent account unless the brain and body are brought back into the formulation.

Consolidation therefore means learning to move from scattered observations to a defensible clinical position. A behavioural change may be real without being diagnostically sufficient. A normal-appearing conversation may coexist with a concerning change in arousal, movement, vision, sensation, continence, coordination, or headache pattern. Conversely, a dramatic symptom report does not by itself establish structural disease. The psychiatrist's work is to observe carefully, state uncertainty accurately, and organise the next step so that neither psychiatric care nor medical safety is abandoned.

17.1 Build a formulation before ordering a test

A **neurological lens** asks a different first question from a checklist: what kind of problem is this patient having now? Begin by describing the clinical state in ordinary, observable terms. Is the dominant disturbance in mood, thought, perception, behaviour, cognition, arousal, movement, or bodily function? Which parts are witnessed, which are reported, and which are inferred from collateral history? This prevents the reflex of treating an investigation as a substitute for formulation.

The useful product is a **syndrome formulation**: a concise account linking the mental-state change to its onset, progression, associated bodily symptoms, treatment exposure, background vulnerability, and consequences for safety. It should be able to accommodate competing explanations. A patient can have a primary psychiatric disorder and a concurrent neurological illness; the choice is not invariably one or the other. Equally, an unusual psychiatric presentation need not be made neurological merely because it is difficult to classify. Formulation keeps the question open while making the next observation purposeful.

Ask what would be surprising if the working explanation were correct. If a presumed affective episode is accompanied by new objective focality, fluctuating arousal, an unfamiliar neurological complaint, or a striking departure from the person's usual illness pattern, the formulation needs

revision. If a presumed neurological disorder is being considered solely because symptoms are intense, the evidential basis may still be weak. The essential habit is to make the discrepancy explicit rather than silently absorbing it into a psychiatric label.

RULE **Investigations answer a formulated question; they do not create one.** State the clinical concern, the competing explanation, and the decision that a result could change before requesting the test.

17.2 Establish the state and the trajectory

The most transferable neurological skill in psychiatry is attention to state. A **state-dependent assessment** recognises that a mental examination is valid only to the extent that the patient can sustain the level of arousal, attention, communication, and cooperation required for it. Confident interpretation of thought content is hazardous when the underlying state is unstable. In that setting, the task shifts from refining psychopathology to describing the disturbance, protecting the patient, and clarifying its cause.

Trajectory gives the state meaning. Obtain an account of what the patient was like before the current change, how the change became apparent, whether it has been stable or variable, and what else changed at the same time. Collateral history is often the bridge between a convincing bedside interview and a misleading formulation. It may reveal a loss of previously routine skills, altered sleep-wake behaviour, a change in gait, episodes of unresponsiveness, new falls, altered intake, or a treatment-related temporal relationship that the patient cannot report reliably.

Do not let the word “chronic” end the enquiry. A longstanding psychiatric diagnosis does not explain every later change, and a longstanding neurological symptom does not make a new deterioration benign. The clinically important question is whether the present pattern is continuous with the patient's established baseline. When it is not, record the difference in concrete language. This keeps the handover useful and makes later reassessment possible.

- Describe the current level of engagement before interpreting elaborate mental content.
- Seek collateral evidence of functional change rather than relying only on diagnostic history.
- Record the sequence of symptoms and treatment exposures in the patient's own clinical timeline.
- Reassess the formulation when the course changes rather than defending the original label.

GOOD TO REMEMBER

A careful description of the change from baseline is often more clinically valuable than a premature conclusion about its cause. It allows the next clinician to see whether the patient is improving, fluctuating, or evolving in a new direction.

17.3 Give positive and negative findings their proper weight

A **positive neurological finding** is an observed abnormality that can be described, repeated, and placed in context. It gains value when it is internally consistent, corroborated by another source of information, or changes the working formulation. A **negative neurological finding** is not a

declaration that neurology has been excluded. It is a statement about what was looked for, under what conditions, and what was not found at that time. This distinction is vital in psychiatric settings, where agitation, withdrawal, fear, pain, medication effects, and limited cooperation can constrain the examination.

Use the examination already taught in this chapter as a targeted extension of the mental-state assessment. Observe spontaneous movement before asking the patient to perform tasks. Notice facial animation, speech effort, posture, balance during ordinary activity, and the way a patient enters or leaves the room. Then examine the relevant domains deliberately, explaining what could and could not be assessed. The full techniques belong to the physical neurological examination section; the consolidating skill is deciding which finding is meaningful for this patient.

Negative findings matter because they limit overinterpretation. They can support a plan to monitor rather than escalate, or they can identify that the concern remains unexplained despite an unrevealing bedside screen. Neither use is passive. The note should distinguish “not present” from “not assessed” and “not assessable.” That small discipline protects against the false reassurance created by a templated normal examination.

TRAP **Calling an examination normal when it was incomplete.** In psychiatric practice, non-cooperation or altered arousal is itself clinically relevant; document the limitation and decide whether observation, collateral information, or a different setting is needed.

17.4 Localise cautiously and think in patterns

Localisation is a reasoning tool, not a performance of anatomical recall. It asks whether a cluster of findings points towards a coherent level or system, whether the pattern is diffuse rather than focal, or whether the findings do not yet support either conclusion. Its value lies in directing attention and communication. It should never be used to make a lesion diagnosis from a single sign, especially when the sign may be influenced by anxiety, medication, effort, pain, or the examination context.

Think in patterns across domains. A change in behaviour becomes more concerning when it travels with a change in consciousness, cognition, motor function, sensory experience, autonomic function, or cranial symptoms. By contrast, a symptom that remains isolated, stable, and compatible with the established clinical picture may call for follow-up rather than an urgent neurological narrative. Pattern recognition must be paired with the humility to say that the pattern is incomplete. The detailed anatomical framework belongs in the Functional Neuroanatomy and Neurophysiology notes; the psychiatrist needs enough applied reasoning to notice when the phenotype does not fit.

Be especially careful with apparent inconsistency. Inconsistency can arise from fluctuating disease, fatigue, communication difficulty, distress, changing attention, pain, medication effects, or an examination technique that has not controlled the task. It may be diagnostically informative, but it is not a shortcut to a conclusion. Repeat the observation, compare it with function outside the formal examination, and describe rather than editorialise.

CLINICAL QUESTION	PSYCHIATRIC FORMULATION REMAINS PLAUSIBLE WHEN	NEUROLOGICAL CONCERN STRENGTHENS WHEN	USEFUL NEXT MOVE
Change in mental state	The presentation remains coherent with history and observed state	The state or course is discordant with the working account	Reconstruct baseline and trajectory
Behavioural disturbance	Observed behaviour and collateral information converge	Behaviour coexists with unexplained bodily or cognitive change	Target the bedside examination
Abnormal movement	The phenomenon is described before it is labelled	It is new, progressive, or accompanied by other signs	Review treatment exposure and examine in context
Reported sensory symptom	The symptom can be followed without premature inference	There is an associated objective or functional pattern	Clarify distribution, consistency, and impact
Uncertain findings	Limitations are recorded honestly	Uncertainty itself affects safety or disposition	Seek senior, medical, or neurological input

17.5 Resist diagnostic overshadowing without abandoning psychiatry

Diagnostic overshadowing occurs when a known psychiatric diagnosis causes clinicians to interpret a new physical or neurological problem as merely another expression of that diagnosis. It is not prevented by indiscriminate testing. It is prevented by returning to the patient's present phenotype, comparing it with baseline, and asking whether the explanation has become too convenient. The converse error is also real: treating all distress, bodily vigilance, or atypical expression as evidence of neurological disease can disrupt appropriate psychiatric care.

The balanced approach is dual-track reasoning. Continue to address immediate psychiatric risks, distress, communication needs, and treatment adherence while pursuing a proportionate medical or neurological assessment. This avoids the harmful pause in care that follows a false choice between "psychiatric" and "organic." It also makes the consultation more respectful. Patients and families should not have to prove that physical symptoms deserve attention because a psychiatric diagnosis already exists.

Medication review is part of this lens. Ask whether the observed change predates treatment, follows an exposure, changes with adherence, or has been previously documented. Do not infer causality from temporal proximity alone. Compare the suspected adverse effect with the whole clinical picture and discuss uncertainty clearly. The detailed adverse-effect framework is covered in the Schizophrenia: management, antipsychotics and adverse effects notes and the Mood disorders: pharmacotherapy notes.

PITFALL

Writing “behavioural” as though it were an explanation. Behaviour is a description of what the patient is doing. A formulation still has to account for the state, trajectory, context, associated signs, and immediate risks.

17.6 Match the investigation to the question

Investigations become clinically useful when their limitations are anticipated. Neuroimaging, electrophysiology, and cerebrospinal-fluid examination answer different questions and occupy different places in the work-up. The psychiatrist should be able to state the concern that prompted the request, identify what a positive or negative result would mean for the current plan, and recognise when an apparently reassuring result does not settle the clinical question. Technical details of imaging belong in the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes; the psych-work-up role of electrophysiology belongs in the Epilepsy and psychiatry notes; cerebrospinal-fluid pathways require the Neuropsychiatry of systemic and neurological disease notes where relevant.

Before requesting an investigation, clarify the clinical decision at stake. Are you seeking evidence to explain a change in consciousness, a seizure-like episode, a focal complaint, a progressive cognitive change, an inflammatory or infectious process, or an otherwise unexplained neuropsychiatric state? The wording should make the question visible to the receiving team. “Rule out organicity” is not a useful request because it names neither a syndrome nor an action.

After a result returns, bring it back to the bedside. A test result is not self-interpreting, and an incidental finding is not automatically causal. Reconcile it with the state examination, temporal pattern, and prior probability established by the history. If the result does not fit, do not silently select the more convenient story. Clarify the discrepancy with the appropriate specialist and update the written formulation.

- Name the symptom cluster or observed change that generated the request.
- State what decision a result could alter: disposition, monitoring, treatment, or referral.
- Record limitations in timing, state, and examination quality.
- Interpret the result alongside the patient, not instead of the patient.

17.7 Apply the lens across common psychiatric interfaces

The same reasoning method applies whether the presenting interface is headache, sleep, altered consciousness, movement, cognition, or a possible seizure-related event. The content differs, but the sequence does not: establish the state, recover the trajectory, examine for discriminating signs, identify the safety question, and select the appropriate pathway. This chapter supplies the neurological frame; it does not replace the disorder-specific teaching.

For a patient with headache and psychiatric symptoms, ask whether the headache pattern, associated neurological complaints, visual experience, and treatment context require a broader account than stress or somatisation. The disorder-specific discussion belongs in the headache sections of this chapter. For sleep complaints, do not reduce daytime behaviour or cognitive

symptoms to insomnia without considering the sleep-wake context; the disease-level assessment and management are in the Eating and sleep-wake disorders notes. For altered consciousness, distinguish descriptive psychopathology from the neurological problem of arousal and responsiveness, and use the Stroke, TBI, delirium and focal brain syndromes notes for delirium in full.

Movement complaints demand particular care because language is imprecise. Patients may use the same words for inner restlessness, pain, tremor, weakness, cramps, or an urge to move. Begin with the patient's experience and observed behaviour, then review treatment exposure and perform the relevant examination. Do not collapse these possibilities into a single psychiatric label. The distinction between restless legs phenomena and akathisia is addressed in the sleep and movement material, while adverse-effect teaching belongs in the Schizophrenia: management, antipsychotics and adverse effects notes.

STATE	TIME	SIGNS	SAFETY
WHAT IS HAPPENING NOW?	HOW DID IT EVOLVE?	WHAT IS OBSERVED?	WHAT ACTION CANNOT WAIT?

17.8 Communicate uncertainty as an active clinical plan

Diagnostic humility is not vagueness. It is the ability to state what is known, what remains uncertain, what has been examined, and what will trigger a change in plan. This is particularly important when the patient is transferred between emergency, medical, neurological, and psychiatric settings. A useful note permits another clinician to reconstruct the reasoning without guessing which facts were observed and which were assumed.

A good handover names the current psychiatric risks and the concurrent neurological concern in parallel. It specifies the baseline, recent trajectory, pertinent positive and negative findings, limitations of the assessment, investigations requested or reviewed, and the reason for escalation or observation. It also tells the receiving team what deterioration would matter. This is **safety netting**: not an afterthought, but a shared plan for detecting change before it is normalised.

Uncertainty should alter the plan, not merely soften the language. It may justify closer observation, a repeat examination when the patient is more assessable, direct discussion with a medical colleague, or a clearer threshold for transfer. The appropriate response depends on the setting and the immediate risk, but the reasoning can be communicated consistently: the available evidence does not yet account adequately for the presentation, so the team is actively preserving safety while seeking clarification. That is a clinically strong position, not a failure of diagnostic confidence.

Clinical confidence should rise only when the account becomes more coherent. When it does not, seek help early and preserve the evidence that prompted concern. The psychiatrist is not expected to provide the final answer to every neurological problem. The psychiatrist is expected to notice when the current answer does not adequately explain the patient and to act in a way that keeps the patient safe while expertise is brought in.

MNEMONIC

SCOPE

State: define the patient's current level of function and arousal. **C**ourse: reconstruct the change from baseline. **O**bservations: separate observed signs from reported symptoms and untested assumptions. **P**attern: ask whether the findings cohere with the working formulation. **E**scalation: state the safety question and the next responsible service or setting.

When the presentation does not cohere, return to state, course, signs, pattern, and safety before deciding that it is "just psychiatric."

References

The canonical sources this chapter is grounded in. Every fact-bearing specific in the notes traces to one of these; where a primary paper is named in the text, it is cited inline. The list is ordered by how much of the chapter rests on each.

1. Kaufman DM, Geyer HL, Milstein MJ. Kaufman's Clinical Neurology for Psychiatrists. 9th ed. Elsevier.
2. Boland R, Verduin ML, eds. Kaplan and Sadock's Comprehensive Textbook of Psychiatry. 11th ed. Wolters Kluwer; 2024.
3. Brazis PW, Masdeu JC, Biller J. Localization in Clinical Neurology. 8th ed. Wolters Kluwer.
4. Glynn M, Drake WM, eds. Hutchison's Clinical Methods. Elsevier.
5. Tripathi KD. Essentials of Medical Pharmacology. 8th ed. Jaypee Brothers.
6. Clinical Methods in Psychiatry. 1st ed. All India Institute of Medical Sciences; 2024.
7. Nelson JC, Meyer JM, Schatzberg AF, et al. The Lithium Handbook. Cambridge University Press.
8. American Psychiatric Association. Diagnostic and Statistical Manual of Mental Disorders, 5th ed, Text Revision (DSM-5-TR). APA; 2022.
9. Geddes JR, Andreasen NC, Goodwin GM, eds. New Oxford Textbook of Psychiatry. 3rd ed. Oxford University Press.
10. Tasman A, Riba MB, Alarcon RD, et al, eds. Tasman's Psychiatry. 5th ed. Springer.
11. Taylor DM, Barnes TRE, Young AH. The Maudsley Prescribing Guidelines in Psychiatry. 15th ed. Wiley-Blackwell.
12. Stahl SM. Stahl's Essential Psychopharmacology: Neuroscientific Basis and Practical Applications. 5th ed. Cambridge University Press.
13. Cowen P, Harrison P, Burns T. Shorter Oxford Textbook of Psychiatry. Oxford University Press.
14. Martin A, Bloch MH, Volkmar FR, eds. Lewis's Child and Adolescent Psychiatry. 5th ed. Wolters Kluwer.
15. Casey P, Kelly B. Fish's Clinical Psychopathology. 4th ed. Cambridge University Press.
16. Brunton LL, Knollmann BC, eds. Goodman and Gilman's The Pharmacological Basis of Therapeutics. McGraw-Hill.
17. National Institute for Health and Care Excellence. Headaches in over 12s: diagnosis and management. NICE clinical guideline CG150.
18. Teasdale G, Jennett B. Assessment of coma and impaired consciousness: a practical scale. Lancet 1974;2(7872):81-84. The Glasgow Coma Scale is the copyright of Glasgow University and is reproduced here with that acknowledgement.
19. Buchanan RW, Heinrichs DW. The Neurological Evaluation Scale (NES): a structured instrument for the assessment of neurological signs in schizophrenia. Psychiatry Res 1989;27(3):335-350. Cited as an instrument; no item content is reproduced.
20. Transplantation of Human Organs and Tissues Act and Rules, Government of India, as amended.

Additional sources cited inline: Cicchetti D, ed. Developmental Psychopathology; Jiloha RC. Forensic Psychiatry: An Indian Perspective; Schizoobsessive Disorder (Cambridge University Press). Guideline positions are attributed in the text to the body

that issues them, and where a guideline position and Indian practice diverge, both are named and neither is presented as the other.

PART III

Dementias and Neurocognitive Disorders

Telling the dementias apart while there is still something to do about it, and treating the behaviour without sedating the person.

01 The major dementias, from Alzheimer's to the frontotemporal syndromes

02 The treatable, the rapid and the borderline

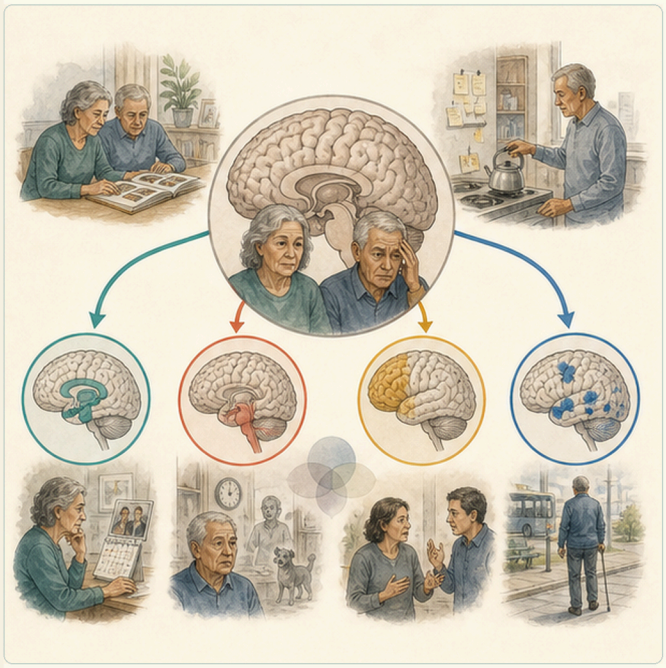
03 Assessment, biomarkers and the pharmacology of dementia

04 Consolidate and apply

05 References

VISUAL 3.1

Four major dementia phenotypes



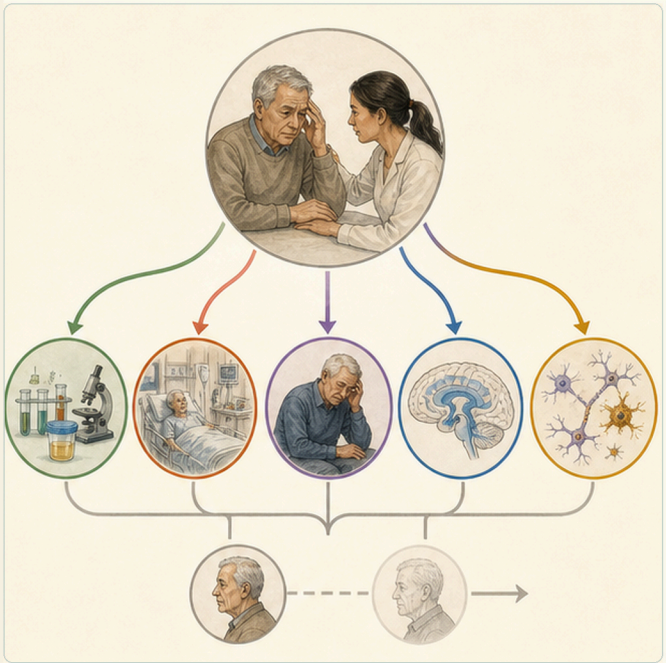
- Alzheimer pattern
- Lewy-body pattern
- Frontotemporal pattern
- Vascular pattern
- Mixed disease

Major dementias are compared through the leading cognitive-behavioural pattern, neurological clues and course while mixed pathology remains possible.

Source lock: day-32.final.html#sec-t13

VISUAL 3.2

The treatable, rapid and borderline differential

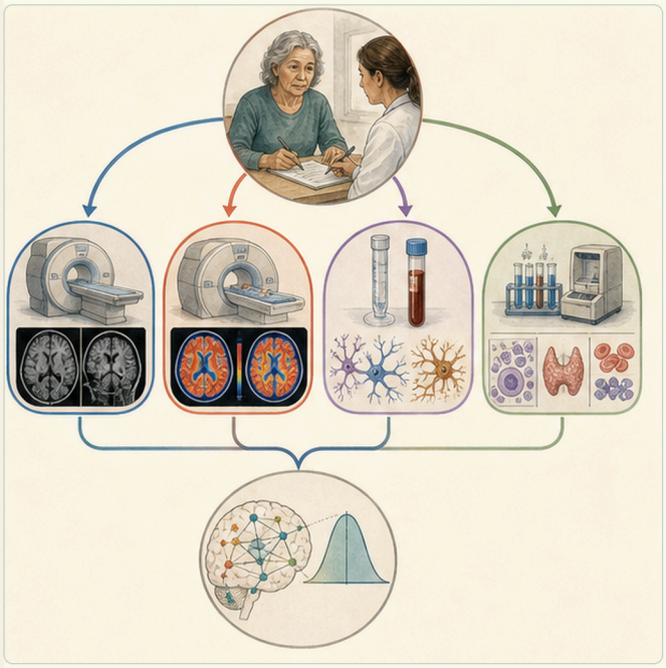


- Reversible causes
- Delirium
- Depression
- Normal-pressure hydrocephalus
- Rapid progression
- Mild cognitive impairment

Cognitive decline is triaged by time course, reversibility, neurological pattern and functional change rather than by a single screening score.

Source lock: day-32.final.html#sec-t19

VISUAL 3.3



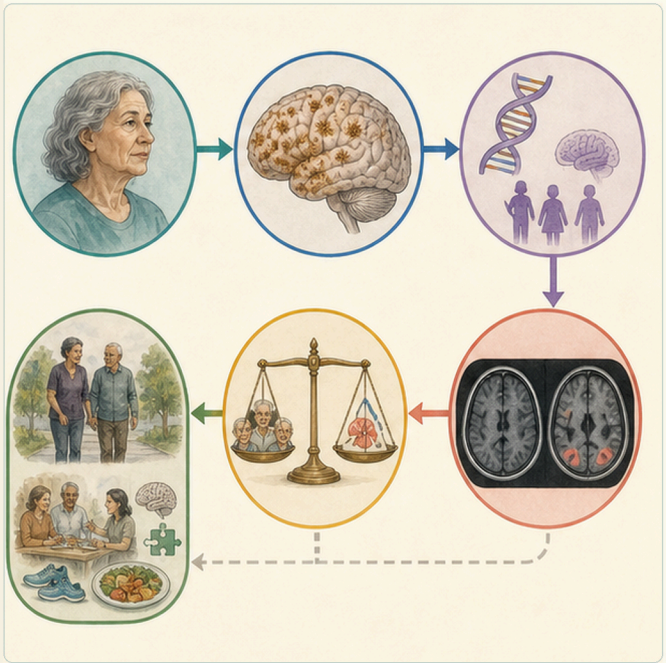
From clinical phenotype to biomarker pattern

- Clinical function
- Structural MRI
- PET
- CSF and plasma
- Reversible laboratory causes

Dementia biomarkers refine a clinical syndrome; they do not replace history, function, neurological examination or the search for treatable causes.

Source lock: day-32.final.html#sec-t22

VISUAL 3.4



Anti-amyloid treatment is a gated pathway

- Early symptomatic stage
- Confirmed amyloid
- APOE and risk context
- ARIA monitoring
- Benefit-harm review
- Supportive care

Anti-amyloid treatment requires stage confirmation, amyloid evidence, risk discussion and MRI safety monitoring alongside continuing supportive care.

Source lock: day-32.final.html#sec-t26

PAPER IV

Dementias and neurocognitive disorders

Three bands . one complete notes document

The major dementias, from Alzheimer's to the frontotemporal syndromes

1 · Dementia for the psychiatrist: what this chapter owns, and what it points to

The marks are earned by organising uncertainty, not by unloading a catalogue of causes. Cognitive decline reaches psychiatry through many doors: forgetfulness, altered judgement, language difficulty, behavioural change, loss of occupational competence, caregiver distress, an abnormal investigation, or a crisis of safety. These presentations do not arrive already sorted into syndrome, cause and consequence. The clinician must construct that order.

The organising task is to keep several questions separate while allowing their answers to inform each other. Is there convincing change from the person's previous level? What kind of cognitive and behavioural pattern is present? How has it unfolded? What has happened to everyday function? Which process best explains the pattern, and what else may be amplifying it? What evidence would alter confidence? What action is required now? The chapter's **diagnostic architecture** is built around these questions. Disease-specific sections supply the criteria, investigations and treatments; this orientation supplies the map into which those details belong.

1.1 Begin with levels of statement

A cognitive complaint, a demonstrated cognitive syndrome and an etiological diagnosis are not interchangeable statements. A complaint describes an experience or observation. A syndrome statement integrates evidence of change, the affected domains, behaviour and functional effect. An etiological statement proposes the disease process or combination of processes responsible. A safe answer makes clear which level is established, which is inferred and which remains uncertain.

This separation matters because the evidence needed at each level differs. A patient may report severe difficulty without showing the expected pattern on assessment. Another patient may deny change despite persuasive collateral evidence and altered performance. A striking investigation may support a proposed pathology but cannot, by itself, describe the person's functional state. Conversely, a convincing syndrome does not automatically identify its biological cause. The formal definitions and classification of neurocognitive disorders are taught in their dedicated section; here the essential orientation is that labels must not outrun evidence.

Complaint	Syndrome	Phenotype	Cause
THE ENTRY	THE LEVEL	THE PATTERN	THE INFERENCE

The idea of **syndrome level** protects both examination writing and practice. It prevents a symptom such as forgetting names from being treated as a disorder, and it prevents a broad term such as dementia from being offered when the question asks for a cause. It also allows

proportionate uncertainty. The clinician can be confident that meaningful decline is present while remaining appropriately cautious about its exact origin. That is not diagnostic weakness. It is accurate calibration.

RULE **State complaint, syndrome, phenotype and cause at the level the evidence can support.** Each is a different claim. Strong reasoning links them explicitly instead of collapsing them into a premature label.

1.2 Read every case through intersecting axes

A long differential becomes manageable when the presentation is placed on a small set of clinical axes. Baseline asks what the person could previously do. Pattern asks which cognitive, behavioural and neurological functions have changed together. Course asks how the change emerged and evolved. Function asks what the person can still manage independently and what others have quietly taken over. Context asks which medical, psychiatric, medication-related, sensory and social factors distort performance. Safety asks what may cause harm before diagnostic certainty is complete.

The **cognitive phenotype** is the recognisable configuration produced by relative deficits, preserved capacities, associated behaviour and neurological signs. It is more informative than the statement that cognition is "poor." The detailed profiles of the major dementia syndromes belong in their disease sections. At the orientation level, remember that similar overall impairment can arise from different patterns, and similar complaints can reflect different mechanisms. Phenotype guides the differential; it does not replace etiological proof.

CLINICAL AXIS	ORIENTING QUESTION	WHAT IT PREVENTS	WHERE DEPTH FOLLOWS
Baseline	What has changed from the person's established abilities?	Equating low performance with acquired decline	Clinical approach and young-onset presentations
Pattern	Which cognitive, behavioural and neurological features travel together?	Assuming every dementia is primarily amnestic	Major disease-family sections
Course	How did the change begin, fluctuate and progress?	Ignoring urgency or superimposed illness	Treatable and rapidly evolving presentations
Function	Which real-world tasks have changed, and why?	Letting a test score substitute for daily-life evidence	Definitions, assessment and care planning
Context	What could mimic, worsen or conceal the syndrome?	False either-or thinking about psychiatric and neurological causes	Systematic workup and differential diagnosis
Evidence	Which findings support, weaken or redirect the formulation?	Allowing a scan or biomarker to become the diagnosis	Assessment, imaging and fluid biomarkers
Safety	What requires action even while uncertainty remains?	Postponing care until every label is settled	Behavioural management, caregiver support and future care

Tempo deserves special attention because it changes the meaning of the same symptom. A chronic progressive narrative, a fluctuating presentation, a stepwise course and a rapidly evolving decline open different diagnostic routes. The dedicated sections teach those routes in full. This orientation uses tempo only as a sorting axis: establish it from dated examples and collateral history, then ask whether the proposed disease process genuinely fits it.

1.3 See the chapter as a sequence of decisions

The chapter is easiest to retain when its sections are read as answers to successive clinical decisions. The definitional framework determines what kind of syndrome is being described. The clinical approach establishes how history, collateral information, examination and investigations are assembled. The major disease-family sections then show how distinct phenotypes, pathologies and management questions modify that common approach. The treatable, rapidly evolving and borderline presentations guard against premature closure. Assessment and biomarker sections show how evidence is strengthened. Pharmacological and psychosocial sections translate diagnosis into care.

CHAPTER TERRITORY	QUESTION IT ANSWERS	HOW TO USE IT
Conceptual framework	What counts as a neurocognitive syndrome, and how is severity expressed?	Use it to frame the opening of an answer
Clinical approach	How is change established and patterned?	Use it to structure history, collateral evidence and examination
Major disease families	Which phenotype and pathology best explain the presentation?	Use disease-specific criteria only after the syndrome is clear
Treatable and rapid presentations	What must not be missed or dismissed as inevitable degeneration?	Use tempo and context to determine urgency
Assessment and biomarkers	What evidence can support or challenge the working formulation?	Order and interpret tests as answers to clinical questions
Treatment and care	What should change for the patient and caregiver?	Link each intervention to a target and monitoring plan
Application	Can the knowledge survive a vignette, viva or long answer?	Return to the same reasoning sequence under pressure

This structure prevents fragmented recall: an unranked list of diseases, tests and drugs, or detailed disease knowledge without a decision about what the current case requires. Reading the chapter as a pathway turns stored facts into clinical action. It also prevents overreach. A biomarker section does not make biomarkers the starting point, and a medication section does not make prescribing the whole of care.

GOOD TO REMEMBER

Before revising a disease section, write the question that section answers. After revising it, state what evidence would make you choose that diagnosis over its nearest competitors and what finding would make you reopen the case. This converts descriptive knowledge into usable discrimination without forcing false certainty.

1.4 Keep causes, contributors and neighbours distinct

Cognitive impairment rarely respects the boundaries of a textbook. A patient may have a primary neurodegenerative or vascular process, a psychiatric state that worsens performance, medical illness that reduces reserve, medication effects, sensory barriers and a caregiving environment that magnifies disability. A coherent formulation separates the process thought to generate the syndrome from factors that contribute to its expression. It also remains open to **mixed causation** when independent evidence supports multiple meaningful processes.

Neighbouring disciplines need clean boundaries. Delirium and the broader post-stroke cognitive spectrum are taught in the Stroke, TBI, delirium and focal brain syndromes notes. Endocrine, metabolic and nutritional diseases are taught in the Endocrine and metabolic psychiatry notes. Movement-disorder neuropsychiatry, inherited neurological disease and autoimmune

encephalitis are taught in the Neuropsychiatry of systemic and neurological disease notes. This chapter returns to them only through the dementia lens: when they mimic, contribute to, coexist with or alter the assessment of a cognitive syndrome.

The same principle applies to developmental baseline. Intellectual disability as a developmental condition belongs in the Intellectual disability and autism spectrum disorder notes. A dementia chapter asks a different question: whether there has been acquired decline from that individual's own established level. Low baseline performance and acquired deterioration must not be treated as synonyms. The quality of the longitudinal account becomes especially important when conventional test assumptions fit poorly.

TRAP **Do not write a miscellaneous "reversible causes" list and assume the differential is complete.** A contributor may worsen a dementia without causing it, a treatable disorder may coexist with degeneration, and an apparently familiar dementia label may conceal an urgent superimposed process.

A useful causal statement therefore has layers. It names the leading disease process only when justified, identifies active contributors, acknowledges plausible alternatives, and states what evidence would separate them. The aim is not to place every abnormality into the causal chain. It is to explain the observed pattern with the least distortion while preserving clinically important uncertainty.

1.5 Build evidence by convergence

No isolated source of evidence has universal priority. The patient's account shows subjective experience, goals and fears. Collateral history shows change, compensation and risk in ordinary life. Mental state examination identifies psychiatric phenomena that may affect presentation. Cognitive examination samples domains and error types. Neurological examination identifies associated signs. Laboratory studies investigate selected contributors. Imaging and fluid biomarkers may support or challenge a proposed biological process. Their value comes from convergence around a coherent clinical question.

General administration and scoring of cognitive instruments belong in the Psychometry, Assessment and Descriptive Psychopathology notes. The physics of imaging modalities belongs in the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes. The cerebrospinal-fluid examination procedure belongs in the Neurology foundations for psychiatrists notes. Within this chapter, these tools are considered for their dementia-specific role: what they add to phenotype, differential diagnosis, etiological confidence, treatment selection or follow-up.

Diagnostic discordance is not an inconvenience to be edited out. A history may suggest decline while brief testing appears reassuring. Cognitive performance may appear poor while function and trajectory remain inconsistent with the proposed syndrome. Imaging may show abnormalities that do not match the clinical pattern. A biomarker may increase confidence in a pathology without explaining the whole presentation. Discordance should trigger review of

baseline, language, education, sensory function, effort, context, timing, comorbidity and the possibility of multiple processes.

PITFALL

A referral label can become invisible evidence. Once the notes say “dementia,” every repeated question is read as confirmation and every inconsistency as denial. Reconstruct the history from examples, establish the trajectory independently, and ask what observation would make the inherited label less likely.

Evidence must be interpreted in relation to purpose. Screening, etiological support and monitoring ask different questions. A result can be technically valid yet clinically misapplied. The orienting habit is to state the question before choosing the evidence.

1.6 Place function, safety and caregiving beside cognition

Neurocognitive disorders are lived in households, workplaces and communities, not only in testing rooms. Functional inquiry should therefore move beyond a generic statement that activities are impaired. Ask what changed, whether the task was previously familiar, which cognitive difficulty disrupts it, what physical or social barriers contribute, and who now compensates. Family assistance may preserve outward performance while concealing the extent of support required. Equally, inability caused by physical illness, poverty, restricted opportunity or unfamiliarity should not be mislabelled as cognitive loss.

The **care ecology** includes the patient, caregiver, household, available services, finances, language, transport and the practical distribution of responsibility. It shapes what can be assessed, what can be implemented and what risk can be contained. Caregiver distress is not merely an outcome measure. It may signal unsafe supervision, misunderstanding of behaviour, conflict, sleep loss, financial strain or a plan that has become impossible to sustain. The dedicated caregiving and end-of-life section develops this territory in full.

IN INDIA

Assessment should fit the person's language, literacy, educational opportunity, occupation and household role. Family involvement can provide rich longitudinal evidence and practical support, but “the family will manage” is not a care plan. Clarify who will supervise medication, money, travel, meals and appointments; what happens when that person is unavailable; and whether access, cost or distance makes the proposed investigation or follow-up unrealistic.

Safety reasoning runs alongside diagnosis. The relevant domains include medication errors, financial vulnerability, getting lost, unsafe use of transport or appliances, falls, self-neglect, aggression, exploitation and caregiver collapse. The purpose is not to strip autonomy when risk appears. It is to understand the specific hazard, the person's decision-making ability for the relevant choice, available supports and the least restrictive workable response. Legal particulars should be checked against current applicable law rather than recalled from general textbook memory.

1.7 Let formulation drive treatment

Treatment becomes coherent when it follows the formulation. Begin with **LOCUS**: decide whether the current combination of medical stability, behavioural disturbance, diagnostic urgency, supervision and risk can be managed in the usual setting or needs a more supported setting. Then define **FOCUS**: the immediate symptom, risk, contributor or functional problem that treatment is meant to change, followed by the broader goals of preserving ability, reducing distress, supporting the caregiver and planning ahead. Finally choose **MODUS**: pharmacological, psychological, behavioural, environmental, rehabilitative, social, procedural and palliative approaches as appropriate.

The pharmacological sections teach cognitive enhancers, behavioural prescribing and disease-modifying therapy with their dementia-specific indications and safeguards. General antipsychotic pharmacology belongs in the Schizophrenia: management, antipsychotics and adverse effects notes. This separation matters because a drug cannot be transferred from another indication by familiarity alone. In a dementia answer, every prescription needs a target, an expected form of benefit, a harm-monitoring plan and a reason to continue, modify or stop.

Non-pharmacological care is not an appendix added after medicines. Communication, environmental adaptation, meaningful activity, sensory correction, sleep routine, caregiver education, behavioural analysis, rehabilitation and practical support may determine whether a plan succeeds. Use the same logic as for medication: identify the target, choose a feasible intervention and review the outcome.

- Link each intervention to a stated clinical or functional target rather than to the diagnosis in the abstract.
- Separate treatment of the underlying disease process from treatment of contributors, complications and caregiver strain.
- Make feasibility part of efficacy by considering supervision, cost, travel, language and follow-up capacity.
- Review benefit and harm in terms meaningful to the patient and caregiver, not only through a test score.
- Revise the plan when function, behaviour, medical state, support or goals change.

Current disease-specific guidance should be consulted in the relevant management sections rather than reconstructed from memory. An orientation cannot safely substitute for formal eligibility criteria, dosing, monitoring schedules or statutory requirements. Its task is to show where those particulars belong and how they serve a wider clinical plan.

1.8 Use the master map under examination pressure

When the vignette becomes crowded, frame it before you solve it. The mnemonic below is not a diagnostic criterion and does not replace disease-specific knowledge. It is a retrieval scaffold that preserves the order of reasoning. Each element points to material developed elsewhere in the

chapter, allowing the candidate to retrieve detail without turning the answer into an unstructured list.

MNEMONIC

FRAME

Function and former baseline: establish meaningful change in real life. **R**ate and route: describe onset, tempo and trajectory. **A**ffected pattern: identify the cognitive, behavioural and neurological phenotype. **M**imics, modifiers and mixtures: separate cause from contributors and neighbours. **E**vidence and evolving care: integrate collateral history, examination and targeted tests, then connect formulation to LOCUS, FOCUS, MODUS and longitudinal review.

In a long answer, open by stating the clinical construct in plain professional language. Classify only to the degree required, then move through presentation, course, functional change, associated features, examination, targeted investigations, differential diagnosis and management. In a vignette, begin with the most discriminating evidence rather than repeating every fact in the stem. In a viva, make uncertainty audible: state the leading formulation, the evidence supporting it, the strongest alternative and what would resolve the difference.

The chapter should leave the reader with a **longitudinal formulation**, not a static label. Diagnosis may become more or less secure as trajectory, function, examination, treatment response and biomarkers accumulate. Care goals also change. The clinician's responsibility is to revise openly without losing continuity: retain what remains supported, discard what no longer fits, and explain the practical consequence of the revision to the patient and caregiver.

The boundary map remains important during revision. General psychometry, imaging physics, cerebrospinal-fluid procedure, antipsychotic pharmacology, intellectual disability, delirium, systemic and neurological neuropsychiatry, and endocrine or metabolic disease retain their own chapter homes. Cross-reference them when needed, but keep the dementia answer centred on acquired cognitive decline, its phenotype, cause, functional consequences and care.

Frame every case by former function, rate of change, affected pattern, mimics and modifiers, converging evidence, and the care consequences of the formulation.

2 · Dementia: definition, classification and the neurocognitive-disorder framework

The diagnostic centre of gravity is decline plus disability. Dementia is not simply poor memory, a low cognitive-screen score, or any cognitive complaint in later life. It is an acquired syndrome in which decline from a previous level of ability is established, the affected cognitive domains are identified, and the consequences for everyday functioning are judged. The examination marks lie in stating that architecture clearly, then showing that the modern manuals draw its boundaries differently.

The contemporary framework also prevents a common conceptual error: the syndrome is not an aetiology. After establishing severity, the clinician must still specify the likely cause and recognise

that the cognitive phenotype may be cortical, subcortical, or mixed. This section therefore moves from definition to diagnostic threshold, then to practical classification. Disease-specific criteria, investigations and management belong to their respective sections.

DECLINE	DOMAIN	FUNCTION	CAUSE
FROM BASELINE	PROFILE	INDEPENDENCE	AETIOLOGY

2.1 From dementia to the neurocognitive-disorder framework

The older term dementia remains clinically familiar and is still used in the international classification. In the **DSM-5-TR major and mild neurocognitive disorder framework**, major neurocognitive disorder maps broadly onto dementia, while mild neurocognitive disorder maps broadly onto mild cognitive impairment. “Broadly” matters. These are conceptual correspondences, not permission to treat every historical use of the older labels as diagnostically identical.

The umbrella term has a wider scope than dementia. It can describe acquired cognitive impairment across the lifespan, including impairment associated with brain injury or infection, rather than implying a disorder restricted to later life. This wording also directs attention away from memory alone. A person may present because of executive, language, perceptual-motor, attentional, or social-cognitive decline, and the syndrome can still sit within this framework.

RULE **Define the syndrome before naming the disease.** Establish acquired decline, map the cognitive domains, determine whether independence is lost, exclude a delirium-only presentation and a better mental-disorder explanation, and only then add the aetiological formulation.

“Dementia” is therefore best retained as a clinically useful syndromic term rather than used as shorthand for irreversible degeneration. Some causes are degenerative, but others are vascular, neurological, endocrine, nutritional, infectious, metabolic, traumatic, or exposure-related. The classification should keep syndrome, severity and cause on separate lines because each answers a different question.

2.2 Major neurocognitive disorder: the diagnostic architecture

The DSM formulation begins with evidence of significant cognitive decline from the person's previous level in **one or more** cognitive domains. Evidence has a clinical component and an objective component: concern may come from the person, a knowledgeable informant, or the clinician, while substantial impairment should preferably be documented by standardised neuropsychological testing or, when that is unavailable, another quantified clinical assessment. Neither component should be treated as decorative. A test score without a credible decline history may misclassify limited education, language mismatch or longstanding low ability; a complaint without objective corroboration may reflect another process.

The named domain map contains **six** areas: complex attention, executive function, learning and memory, language, perceptual-motor function, and social cognition. The list is diagnostically liberating because memory impairment is not mandatory in every phenotype. It is also clinically

demanding: a routine interview that samples only orientation and recall can miss an executive, visuospatial, language-led or social-cognitive syndrome.

MNEMONIC

Decline - domains - daily function - differential

Use this sequence when constructing the definition. First establish decline from baseline; then identify the affected domains; next decide what has happened to independent daily function; finally check delirium and alternative mental-disorder explanations.

The decisive next question is whether cognitive deficits interfere with independence in everyday activities. At the **major functional threshold**, the person needs assistance with complex instrumental activities such as managing bills or medications. The criterion is about capacity, not merely whether another person happens to perform a task. A family may have always managed finances, or may take over tasks prematurely; conversely, a person may remain technically independent only because routines have narrowed and risks are being concealed. Functional enquiry must reconstruct what the person previously did, what now fails, how much support is required, and whether the failure is attributable to cognition.

The remaining **delirium and mental-disorder exclusions** are equally important. The deficits must not occur exclusively during delirium and must not be better explained by another mental disorder. This does not mean dementia and delirium cannot coexist, or that a patient with a mood or psychotic disorder cannot also develop cognitive disease. It means that the clinician must not assign the cognitive syndrome when the entire presentation is sufficiently explained by the competing condition.

GOOD TO REMEMBER

Ask for examples of failed tasks rather than asking only, “Is the patient independent?” Medication errors, missed payments, unsafe travel and abandoned household sequences reveal cognitive disability more reliably than a global yes-or-no answer.

2.3 Mild versus major: function is the hinge

Mild neurocognitive disorder also requires decline from a previous level within the same domain set, supported by concern and modest objective impairment. The separating feature is that the cognitive deficits do not remove the capacity for independence. Complex instrumental activities remain preserved, although they may require greater effort, compensatory strategies or accommodation. Thus “mild” refers to functional impact within this diagnostic framework, not to how distressing the symptoms feel, how abnormal a score looks, or how serious the underlying disease may eventually become.

AXIS	MILD NEUROCOGNITIVE DISORDER	MAJOR NEUROCOGNITIVE DISORDER
Magnitude of decline	Modest decline with modest objective impairment	Significant decline with substantial objective impairment
Cognitive breadth	Any named cognitive domain may be affected	Any named cognitive domain may be affected
Independence	Capacity remains intact	Deficits interfere with independence
Complex instrumental tasks	Preserved, but may need extra effort, strategy or accommodation	Assistance is required at least for complex instrumental tasks
Core clinical question	How is independence being maintained?	Which cognitive failures now require another person's help?

Severity within major neurocognitive disorder is also functionally specified. Mild severity means difficulty with instrumental activities such as housework or money management; moderate severity means difficulty with basic activities such as feeding or dressing; severe illness means full dependence **functional severity specifiers**. Do not confuse "mild major neurocognitive disorder" with mild neurocognitive disorder. The former has crossed the loss-of-independence threshold; the latter has not.

PITFALL

Do not equate family supervision with demonstrated dependence. Clarify whether help is necessary because cognition has failed, merely convenient, culturally customary, or prompted by physical disability. Functional staging is invalid if the cause of assistance is not established.

2.4 ICD framework and the cross-system examination trap

The international classification uses a similar clinical logic but a different domain threshold. Under **ICD-11**, dementia is characterised by marked impairment in **two or more** cognitive domains, severe enough to cause significant impairment in important areas of functioning. Mild neurocognitive disorder may involve similar domains, but its deficits are not severe enough to cause significant functional impairment. Memory is affected in most forms of dementia, yet impairment is not restricted to memory; executive functioning, attention, language, social cognition and judgement, psychomotor speed, and visuoperceptual or visuospatial functioning may be involved.

TRAP **Do not harmonise the domain count across manuals.** The DSM criterion uses the lower domain threshold stated above, whereas the ICD dementia requirement uses the stricter multiple-domain threshold. State the system before stating the threshold.

The international framework separates dementia from normal ageing by magnitude relative to age expectations and by functional impact. Appropriately normed standardised assessment can help establish whether performance deviates from expected ageing. Clinically relevant difficulty

that remains consistent with normal ageing should not be inflated into dementia merely because the patient is worried or slower than before.

The delirium boundary is temporal and phenomenological. Delirium produces a precipitous, transient and fluctuating global neurocognitive disturbance with confusion, whereas dementia is more often gradual, progressive and expressed through impairment of particular cognitive skills. Dementia increases vulnerability to delirium, so a person with established dementia who develops an acute disturbance of attention, orientation and awareness requires an additional delirium diagnosis rather than having the change attributed to "worsening dementia." Delirium itself is taught in the Stroke, TBI, delirium and focal brain syndromes notes.

Severity extension codes are available in the international system and are appended to the aetiological dementia code, including none, mild, moderate and severe categories **ICD severity extensions**. The available source does not establish the functional thresholds defining those levels, so the DSM functional definitions should not be copied into the international system by assumption.

2.5 Classification by cause: a map for differential diagnosis

Once the syndrome has been established, classification by cause prevents the vague and unsafe endpoint of "dementia, unspecified." A useful aetiological map groups causes as neurodegenerative, vascular, other neurological, endocrine, nutritional, infectious, metabolic, traumatic, and exposure-related **aetiological classification**. This is not a claim that every disease within a group behaves alike. It is a search structure that keeps the clinician from considering only the common degenerative disorders.

- **Degenerative and vascular:** consider a primary neurodegenerative process, vascular injury, or a mixed contribution.
- **Neurological and structural:** consider disorders that disrupt brain tissue, white matter, circulation or cerebrospinal-fluid dynamics.
- **Systemic states and acquired insults:** consider endocrine, nutritional, metabolic and infectious causes, together with trauma and relevant exposures.

An alternative map divides causes into inherited and acquired groups. The acquired branch can then be subdivided into primary degenerative, vascular, infective, inflammatory, neoplastic and traumatic causes. These systems are complementary. The first is broad and practical for an initial differential; the second makes inheritance and acquired brain disease explicit. Neither replaces a patient-specific formulation based on tempo, neurological findings, systemic clues, functional trajectory and the cognitive phenotype.

The historical **ICD-10** account remains worth recognising because older questions and records may use it. It retained a syndromic approach, provided general dementia criteria and criteria for **four** specific types, and required decline in activities of daily living. It must be labelled as the earlier classification rather than presented as current international wording.

For disease-level accounts of several categories in the aetiological map, follow the relevant cross-references. The endocrine, metabolic and nutritional diseases themselves are covered in the Endocrine and metabolic psychiatry notes. Inherited and other neurological diseases outside the dementia lens are covered in the Neuropsychiatry of systemic and neurological disease notes. The present purpose is to keep these causes visible in the dementia differential, not to reproduce their disease teaching.

2.6 Cortical and subcortical cognitive phenotypes

A second classification asks where network disruption is most strongly expressed. The distinction between **cortical and subcortical dementia** is anatomical and functional: in the cortical pattern, pathology directly affects cerebral cortical systems; in the subcortical pattern, the major cognitive impact is understood as loss of the normal regulatory influence of subcortical structures on the cortex, particularly prefrontal systems. This distinction is useful because it predicts the pattern of failure more than the global amount of impairment.

Alzheimer disease is the cortical prototype: early disturbance prominently involves episodic memory, language and visuospatial ability, while attention and frontal-executive functions are relatively preserved early. By contrast, **Huntington disease and progressive supranuclear palsy** exemplify the subcortical pattern, in which attentional control and executive function are disproportionately impaired and thinking is slowed.

CLINICAL AXIS	CORTICAL-PREDOMINANT PATTERN	SUBCORTICAL-PREDOMINANT PATTERN
Dominant cognitive failure	Memory, language and visuospatial systems	Attention, executive control and cognitive activation
Speed	Slowing is less prominent early in the prototype	Bradyphrenia is characteristic
Memory mechanism	Severe episodic amnesia may appear early in the prototype	Poor attention and encoding, with a prominent retrieval deficit
Recall versus recognition	Interpret within the broader amnestic pattern	Recognition is typically better than spontaneous recall
Focal cortical signs	Aphasia, apraxia or agnosia may occur	These are relatively absent early
Behaviour and mood	Less prominent early in the prototype	Change in mood, personality and social conduct is common

The memory contrast deserves emphasis. In a subcortical pattern, inefficient attention and encoding combine with impaired retrieval, so the patient may produce little spontaneously yet recognise more when choices or cues are supplied. This differs from the marked early episodic amnesia of the cortical prototype, but it remains a bedside clue rather than a stand-alone aetiological test. Depression, fatigue, education, language, motor slowing and test conditions can also alter apparent retrieval efficiency.

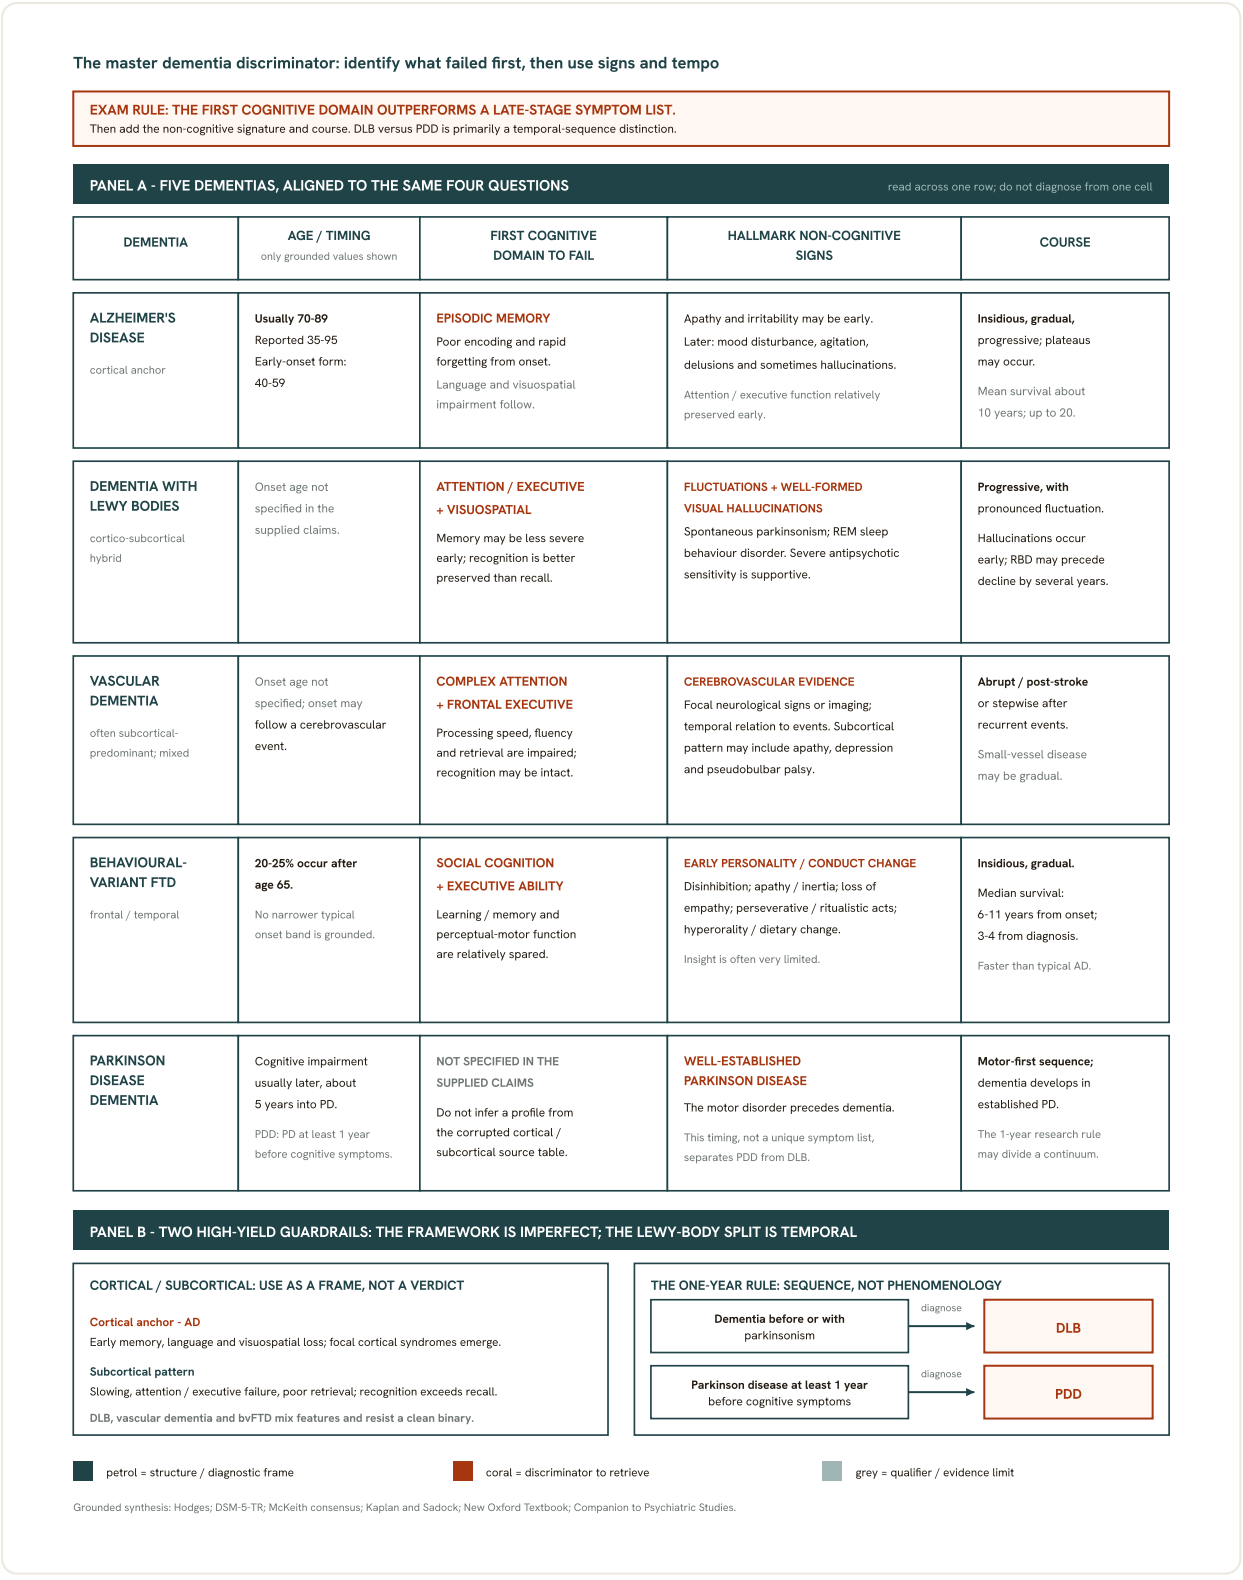


Fig 32.1 Master dementia discriminator: compare age or timing, the first cognitive domain to fail, hallmark non-cognitive signs and course across Alzheimer's disease, dementia with Lewy bodies, vascular dementia, behavioural-variant frontotemporal dementia and Parkinson disease dementia; the cortical/subcortical frame is useful but imperfect, and DLB versus PDD is separated principally by temporal sequence.

The accompanying figure extends phenotype-based discrimination across the major dementia syndromes. Use it after defining the neurocognitive syndrome, not as a shortcut that converts a striking feature into a diagnosis. The profile must be reconciled with onset, course, functional change, neurological examination and the broader differential.

2.7 Why the cortical-subcortical dichotomy is a heuristic

Real disorders do not respect a binary diagram. Vascular dementia often has a subcortical-predominant profile because of lacunar and diffuse white-matter pathology, yet focal cortical damage may coexist. **Dementia with Lewy bodies** combines poor attention and executive ability with cortical dysfunction implicating posterior cortical systems. Behavioural frontotemporal dementia may combine marked personality and social-conduct change with recognisable cortical features. These are mixed cortical-subcortical profiles rather than clean occupants of either pole.

A broader account therefore allows a cortico-subcortical hybrid category. The term is valuable when it stops forced classification; it becomes unhelpful if it is treated as a final diagnosis. Anatomical phenotype and aetiology are different axes. A mixed profile may arise because a disease affects several networks, because pathology has advanced, or because pathological processes coexist.

- Use the dichotomy to organise the bedside cognitive phenotype.
- Do not use it to exclude a disease merely because a feature falls on the other side.
- Reassess the profile over time, while keeping aetiological, functional-severity and anatomical classifications separate in the formulation.

2.8 A formulation that remains clinically useful

A complete formulation should answer a sequence of linked questions. Is there credible acquired decline? Which domains are affected? Is the impairment objectively supported? Has capacity for independent everyday activity been lost? Is the disturbance exclusively delirious or better explained by another mental disorder? Which aetiological class and cognitive phenotype best fit? This sequence turns “memory complaint” into a defensible syndrome statement without pretending that classification alone establishes cause.

Baseline reconstruction is the quiet skill behind all of these decisions. Education, occupation, language, sensory ability, motor impairment, household roles and opportunities for independent activity shape both test performance and the visibility of decline. The clinician should therefore triangulate the person's account, collateral history, observed performance and records of real-world tasks. Disagreement among these sources is information, not noise: the person may lack awareness, an informant may overstate impairment, or a structured clinic encounter may conceal failure that appears only in complex daily life. The aim is not to reward a persuasive narrative but to establish a coherent change from that individual's own baseline.

A useful written statement separates certainty at each level. The cognitive syndrome may be well established while the cause remains provisional; conversely, a known neurological disease does not by itself prove that current functional loss is cognitive. Record the dominant domains, the effect on independence, the temporal course, relevant exclusions and the leading aetiological class. This makes later revision possible without rewriting the history and prevents the disease label from swallowing the syndrome description.

The same sequence also prevents predictable errors. Normal ageing is not dementia without functional and normative evidence. Mild neurocognitive disorder is not simply an early synonym for major disorder. Major disorder is not defined by memory failure alone. A subcortical profile is not synonymous with a particular disease, and a cortical profile is not proof of the cortical prototype. Finally, an acute attentional disturbance in a person with dementia is not automatically progression; superimposed delirium must be considered.

Memorise the hierarchy: establish decline, map domains, judge independence, exclude delirium-only and better mental-disorder explanations, then classify severity, phenotype and aetiology.

3 · The clinical approach to cognitive decline, including young-onset dementia

Why it matters clinically. Cognitive decline is not diagnosed by a low score alone. The clinician must establish that there has been a change from the person's previous level, define the affected cognitive and behavioural phenotype, determine its tempo, and show how it has altered everyday function. The marks lie in turning an imprecise complaint such as "forgetfulness" into a coherent syndrome while keeping psychiatric, neurological, medical, substance-related and contextual explanations in view.

The assessment is therefore an exercise in triangulation. The patient contributes subjective experience, symptoms and goals; an informant contributes longitudinal observation; examination and cognitive assessment test the emerging hypothesis; and functional evidence shows whether the change matters in real life. Discordance between these sources is information, not noise. It may reflect impaired awareness, over- or under-reporting by an informant, language or educational mismatch, an acute superimposed illness, or a phenotype that a brief screen samples poorly.

RULE **Build the syndrome before naming the disease.** First establish change, tempo, cognitive-behavioural pattern and functional consequence; only then rank aetiologies and select investigations.

STORY	PROFILE	FUNCTION	CONTEXT
CHANGE AND TEMPO	DOMAINS AFFECTED	REAL-LIFE CONSEQUENCE	CAUSE AND MODIFIERS

3.1 Begin with change, not with a label

The opening task is to decide what has changed and whose concern brought the patient to assessment. Clarify the first noticed difficulty, the circumstances in which it appeared, and whether the patient, family, employer or another observer recognised it first. Ask for concrete episodes rather than global judgments. "He is forgetful" is weak evidence; a description of repeated missed obligations despite established routines is clinically interpretable. Equally, an isolated error during stress, sleep loss or illness should not be inflated into a progressive syndrome without longitudinal support.

An accurate account cannot be expected from the **patient alone**, because impaired memory and awareness can make a patient-only history incomplete. The core temporal axes are onset, course and duration. Establish whether onset appeared insidious or sudden; whether the course is gradual, stepwise, episodic or fluctuating; whether change is progressive, static or partly reversible; and whether worsening relates to a medical event, medication change or altered environment. These distinctions do not by themselves diagnose a cause, but they reorganise the differential and determine urgency.

Define the **premorbid baseline** with the same care. Schooling, occupation, language, literacy, habitual division of household tasks, sensory limitations and longstanding personality determine what counts as decline. A person who never managed money cannot be judged impaired because a relative now manages it; conversely, loss of a highly practised but cognitively demanding skill may be meaningful even when basic self-care remains intact. The relevant comparison is the patient with their own earlier functioning, not the patient with an imagined average.

GOOD TO REMEMBER

Ask the informant for the earliest unequivocal change and for a recent concrete example. This often separates a longitudinal decline from a vague family impression and anchors every later question to a real baseline.

3.2 Structure the interview to preserve each account

A useful **staged interview** begins with patient and informant together, continues with the informant alone, and ends with the patient being examined alone. The joint phase identifies the presenting concern and allows observation of interaction. The private informant phase permits discussion of decline, risk, behavioural change and caregiver strain without humiliating or provoking the patient. The private patient phase protects confidentiality, allows direct mental-state assessment and reduces cueing or correction by relatives.

This sequence is not a licence to treat the informant as automatically correct. Record who observed each event, how often they see the patient, whether they benefit from a particular account, and whether family conflict may colour the history. When accounts diverge, document the divergence explicitly. A patient may minimise because of impaired awareness or fear of losing independence; an informant may magnify ordinary errors because of anxiety, exhaustion or a sudden change in expectations. The clinician's task is to test each account against examples, chronology and function.

MNEMONIC

STORY BEFORE SCORE

Secure the **story of change**, its **tempo**, the **observed examples**, the resulting **real-life loss** and the patient's **own account** before interpreting a screening result.

Interview technique matters. Use short, concrete questions, allow time for retrieval, check hearing and comprehension, and avoid turning the encounter into a public test of competence.

Ask permission before sensitive questions where possible. Preserve the patient's dignity while still obtaining a candid collateral account. If no informant attends, seek consent to contact someone who knows the patient's previous and current functioning; meanwhile, state clearly that diagnostic confidence is limited by the absent longitudinal history.

3.3 Take a domain-led cognitive history

The history should follow a **domain map** and must not reduce cognition to memory. Probe each domain through everyday examples, then ask when the problem began, how it has evolved, what compensations are used and what consequence followed. A domain-led approach identifies non-amnestic presentations, exposes uneven profiles and prevents a fluent patient from being declared cognitively intact merely because social conversation is preserved.

DOMAIN	TARGETED ENQUIRY	INTERPRETIVE PURPOSE
Attention and concentration	Slips, absent-mindedness, losing track, or failure when distractions compete	Distinguishes sustained inefficiency from a narrowly described memory complaint
Memory	Recent events, retained personal knowledge, remote information, repetition and benefit from reminders	Separates acquisition, storage and retrieval complaints at the level possible in history
Language	Word finding, naming familiar people, comprehension, reading and writing	Detects language-led decline that may distort apparently "global" testing
Executive function	Planning, organising, solving unfamiliar problems, handling medicines or finances and completing multistep tasks	Shows whether complex goal-directed activity has deteriorated
Calculation	Errors in transactions, bills, change or previously familiar arithmetic	Links a cognitive complaint to practical competence
Spatial and perceptual ability	Route finding, recognising faces or objects and navigating familiar surroundings	Identifies impairment that may present as getting lost or misidentification
Praxis	Using familiar tools, dressing, feeding and carrying out learned sequences despite adequate movement	Distinguishes loss of learned action from weakness or poor effort
Social cognition	Judgment, tact, empathy, inhibition, personality and conduct	Captures behavioural presentations that memory-first questioning misses

Do not accept a domain as impaired merely because the patient says "yes" to a broad question. Establish a before-and-after contrast and obtain an example with ecological validity. Ask whether the error is new, whether it recurs, whether prompts help, and whether the task failed because of cognition, mood, motivation, vision, hearing, language, motor disability or unfamiliarity. The

same observed failure can arise through different cognitive routes; history narrows the mechanism without pretending to replace formal assessment.

3.4 Establish function, safety and behavioural change

Functional status converts symptoms into clinical significance. Ask not only whether a task is completed but how it is completed, with what errors, prompting, supervision, compensatory effort and risk. Driving, money management, medication use and activity planning are particularly informative because they require integration of several abilities. Extend the enquiry to self-care, shopping, housework, work, hobbies and social participation. Structured living-skills measures and functional staging scales can support severity assessment, but they do not replace a contextual account.

- For each important activity, document independence, assistance, supervision and complete dependence in plain language.
- Separate cognitive failure from physical incapacity, lack of opportunity, family takeover and culturally assigned roles.
- Ask what has become unsafe, what remains preserved and what support already prevents harm.

Behavioural and psychological change belongs in the initial history, not in an afterthought. Enquire about altered personality, mood, psychotic experiences, apathy, agitation, anxiety, disinhibition, repetitive motor behaviour, sleep, eating and sexual behaviour. The **Neuropsychiatric Inventory** offers a systematic way to capture the recognised disturbance groups, including delusions, hallucinations, agitation, dysphoria, anxiety, apathy, irritability, euphoria, disinhibition, aberrant motor behaviour, night-time disturbance and appetite or eating change. A structured inventory improves coverage; clinical interpretation still requires context, frequency, impact and a search for precipitants.

> **PITFALL**
>
> Do not call preserved washing and dressing “preserved function” and stop. Complex activities often fail earlier, and relatives may conceal decline by quietly taking over finances, medicines, travel or appointments.

3.5 Use background and context as aetiological evidence

A careful **aetiological background** can redirect the entire formulation. Previous affective illness, vascular disease and prescribed or recreational drug exposure may explain, mimic or contribute to cognitive impairment. Relate each potential factor to the timeline rather than compiling an undifferentiated problem list. Ask whether cognition worsened after an event or exposure, improved when it resolved, or followed a parallel course. Association is not proof, but temporal linkage increases or decreases plausibility.

Broaden this into a **contextual history**: relevant medical and neurological illness, cerebrovascular events and risk, head injury, infection risk, psychiatric illness, alcohol and other substance use, family history of dementia or neurological disease, education and occupational or

chemical exposure. Depression deserves particular attention because it can produce cognitive impairment and may coexist with or precede a dementing illness. Its full differentiation from dementia is addressed elsewhere in this chapter; here, the key is to avoid forcing mood and neurocognitive explanations into an either-or choice.

Family history should be constructed across generations where possible. Record the type of impairment, age pattern in qualitative terms, neurological or systemic accompaniments, and whether diagnoses were verified or merely remembered. Do not convert "an elderly relative was forgetful" into evidence of a familial syndrome. Equally, do not dismiss a cluster of early or unusual neurological presentations because relatives used non-medical language. Family structure, adoption, estrangement and premature death can make an apparently negative history uninformative.

The social history is part of measurement. Language, education, literacy, occupation, migration, sensory impairment and familiarity with test materials affect performance. Poverty, isolation, caregiver availability and access to records shape what can be established and what support is feasible. These factors neither prove nor disprove neurocognitive disease. They determine how confidently observed performance can be interpreted and whether apparent loss reflects decline, deprivation of opportunity or mismatch between the assessment and the person.

3.6 Examine and test the hypothesis, not the patient's worth

Bedside cognitive assessment is hypothesis testing. Select tasks that sample the suspected domains, interpret them against language, education, sensory and motor capacity, and note the manner of performance as well as the final score. Error type, strategy, cueing response, perseveration, fatigue and inconsistency may be more informative than a total. Administration and scoring of general cognitive instruments belong in the Psychometry, Assessment and Descriptive Psychopathology notes; disease-specific cognitive profiles are covered in the dedicated assessment section of this chapter.

A teaching vignette illustrates why **collateral history** can outrank a screen when significant aphasia makes commonly used bedside instruments unreliable. The lesson is not that testing is useless in aphasia. It is that a score cannot answer a question the instrument is no longer able to sample validly. Establish premorbid and pre-event function, describe which responses remain interpretable, and seek domain-appropriate assessment rather than converting language failure into a global cognitive diagnosis.

TRAP **"Low cognitive score equals dementia."** A score is evidence within a formulation. It can be lowered by language, education, sensory loss, motor impairment, mood, acute illness or poor engagement, and it can miss a focal or behaviour-led decline.

The examination should also test competing explanations suggested by the history, but a comprehensive neurological examination, neuroimaging principles and fluid examination procedures belong to their respective specialist sections. Refer to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes for imaging modalities and to the Neurology foundations for psychiatrists notes for the general cerebrospinal-fluid

procedure. The present task is to make investigation question-led: every test should address a ranked possibility, urgency or management consequence.

3.7 Young-onset dementia requires a wider diagnostic aperture

Young-onset dementia begins before **65 years**. Alzheimer disease remains the commonest single cause, accounting for **30-40%** of cases in the cited clinical series. **Frontotemporal lobar degeneration** is a close second, while vascular and inherited neurological causes assume greater relative prominence. This differs from later-life dementia, where the usual degenerative and vascular disorders together account for **at least 80%** of cases. The practical message is not that common diseases disappear in younger people. It is that the residual differential is broader, rarer and more often genetically determined or potentially treatable.

All people with a younger-onset presentation require detailed investigation, and the likelihood of an unusual disorder rises as onset becomes younger. “Detailed” means that the diagnostic question must be reopened rather than simply applying a later-life template more intensely. Recheck whether the dominant change is memory, language, behaviour, executive control, movement, seizures, systemic illness or a mixed syndrome; reconstruct the pedigree; verify tempo; and review exposures and prior neurological events. The exact investigation panel belongs to the systematic workup sections and should be chosen from the phenotype rather than ordered as an unexamined bundle.

Red flags include progression beyond **approximately three MMSE points/year**, a family history containing neurological or systemic features outside the common patterns, and distinctive feature combinations. Limb apraxia with myoclonus and alien-limb phenomena suggests a corticobasal pattern; dysarthria, dysphagia and personality change suggest a frontal syndrome with motor-neurone-disease overlap. These combinations are signposts for specialist investigation, not substitutes for diagnosis.

- Escalate urgency when decline is rapid, fluctuating in an unexplained way, or accompanied by new neurological or systemic features.
- Escalate breadth when the phenotype is unusual, the family pattern is striking, or routine explanations do not fit the chronology.
- Escalate genetic care when the history suggests inheritance, because testing affects relatives as well as the patient.

Developmental context also matters. Dementia in children and young adults is rare and is often associated with lysosomal storage disease. In Down syndrome, dementia occurs in **about 50% or more** of affected people and generally emerges later in adult life. Assessment must demonstrate decline from the individual’s established developmental baseline rather than infer dementia from pre-existing disability. Intellectual disability and the developmental syndrome itself are addressed in the Intellectual disability and autism spectrum disorder notes.

Some inherited neurological and metabolic conditions enter the younger-onset differential but are not taught here. Huntington disease belongs in the Neuropsychiatry of systemic and

neurological disease notes. Wilson's disease and inherited metabolic or storage disorders belong in the Endocrine and metabolic psychiatry notes. Naming these routes matters because a generic "dementia screen" is not an adequate endpoint for a younger patient with an atypical phenotype.

3.8 Synthesize the assessment into a defensible formulation

The final product is not a catalogue of abnormalities. It is a concise formulation that states the evidence for decline, the dominant cognitive-behavioural phenotype, tempo, functional effect, associated psychiatric or neurological features, background contributors, safety concerns, and the degree of diagnostic confidence. Separate observed facts from informant report and from inference. State what remains uncertain and what evidence would change the ranking.

A practical synthesis proceeds in sequential passes. First describe the syndrome without an aetiological label: baseline, onset, course, domains, behaviour, function, awareness and associated features. Then rank likely causes and mimics, explain the evidence for and against each, and specify the investigation or follow-up needed to resolve uncertainty.

This structure protects against premature closure. It also allows the clinician to communicate uncertainty honestly without becoming vague. "Progressive cognitive decline" is insufficient if domains and function are unspecified; a confident disease label is unsafe if chronology and collateral evidence do not support it. When acute confusion is possible, the distinction from delirium must be addressed using the Stroke, TBI, delirium and focal brain syndromes notes. Potentially reversible causes require the systematic workup described in the dedicated section of this chapter, not an improvised memory-based list.

> In cognitive decline, the decisive sequence is: prove change, define phenotype and tempo, establish functional consequence, then investigate the cause.

4 · Alzheimer's disease: epidemiology, risk factors and genetics (APOE, APP, presenilins)

The marks lie in separating burden, risk and causation. **Alzheimer dementia** is common in later life, but an epidemiological association is not a diagnosis, a risk allele is not a deterministic mutation, and neuropathological change is not synonymous with clinical dementia. A strong answer therefore moves in a disciplined sequence: define whose prevalence is being quoted, show the steep relation with age, classify modifiable and non-modifiable risks, and then distinguish rare high-penetrance familial disease from the susceptibility architecture of the much larger sporadic group.

The epidemiology also teaches clinical humility. Prevalence depends on population, setting, diagnostic method and calendar period. Age, survival, vascular health, education and social conditions shape the observed burden. Genetics operates within that context rather than outside it. The pathology and amyloid cascade, clinical staging and diagnostic criteria, and treatment are developed in their own sections; here the task is to understand who develops disease, which exposures alter probability, and how the major genes differ in biological and counselling significance.

6.2 million	13.8 million	5.3%	34.6%
US, LIVING WITH AD DEMENTIA IN 2021	US PROJECTION FOR 2060	US, AGE 65 TO 74	US, OLDER THAN 85

4.1 Epidemiological frame: age dominates the curve

The most robust epidemiological pattern is the **age gradient**. In the United States estimate for 2021, prevalence was **5.3% at ages 65 to 74, 13.8% at ages 75 to 84, and 34.6% above age 85**. These are prevalence figures, not the probability that an individual will inevitably develop disease. They reflect accumulated incidence together with survival. They also explain why demographic ageing can enlarge the number of affected people even if age-specific risk does not change.

The same source estimated that 6.2 million United States residents aged 65 years or older were living with Alzheimer dementia in 2021. It projected 13.8 million people older than 65 years by 2060, including 6.7 million older than 85 years. The projection is best read as a service-planning warning: the subgroup with the greatest age-related burden will contribute heavily to future caseloads. It is not an India estimate and should not be transferred across populations as though geography and ascertainment were irrelevant.

DSM-5-TR provides a differently framed set of estimates. In high-income countries, neurocognitive disorder due to Alzheimer's disease rises from 5% to 10% among people aged 60 to 64 to at least 25% thereafter. Its United States estimate for 2016 was 5.4 million people of all ages, around 200,000 with onset before age 65; it also reported disease in 11% of people aged 65 or older and 32% of those aged 85 or older. These figures need not match a later estimate because the year, sampling frame and diagnostic method differ.

EPIDEMIOLOGICAL STATEMENT	WHAT IT LEGITIMATELY SUPPORTS	WHAT IT DOES NOT SUPPORT
Prevalence rises markedly with age	Age is the strongest population-level risk factor	Age alone establishes a diagnosis
A projected case count rises	Future service burden may increase as the population ages	Every individual has a rising deterministic risk
Two sources quote different prevalence figures	Methods, settings and calendar periods must be stated	One estimate can be silently substituted for the other
Clinical prevalence is measured	The number of people meeting a clinical case definition is estimated	Every person with pathological change has dementia
A subgroup has higher observed prevalence	Biology, survival, comorbidity and social exposure deserve analysis	The disparity is an intrinsic property of group identity

IN INDIA

No India-specific prevalence, incidence or national case-count estimate is available in the cited material for this topic. Use the international figures only with their stated population and year. Quoting a United States estimate as Indian prevalence is not approximation; it is a change of denominator, health system and ascertainment context.

4.2 Onset, sex and population disparities

Although Alzheimer disease is strongly age-related, onset is not confined to a single period. The recorded range extends from **age 35 to age 95**. DSM-5-TR states that symptoms usually begin at ages 70 through 89, while forms beginning at ages 40 through 59 are often, but not invariably, associated with known causative mutations. The useful inference is probabilistic: an unusually early presentation strengthens the reason to examine inheritance carefully, but age at onset by itself neither proves nor excludes a pathogenic mutation.

Observed prevalence is higher among women in the cited United States data: 12% of women older than 65 compared with 6% of men older than 65. The major explanation offered is women's longer lifespan, with possible biological and environmental or cultural contributions. The examination error is to convert prevalence difference into a simplistic sex-specific cause. Longer survival increases the time spent in the age bands where prevalence is greatest, and the remaining difference may reflect several interacting mechanisms.

The same caution applies to racial and ethnic disparities. Among United States residents older than 65, prevalence was 18.6% in Black Americans, 14% in Hispanic Americans and 10% in White Americans. The source attributes much of this difference to unequal burdens of cardiovascular disease and diabetes and to socioeconomic inequities. Group labels are therefore signals to investigate accumulated exposure, access, comorbidity and ascertainment, not biological shortcuts. A doctor-facing account should describe disparity without racial essentialism.

Depending on setting and diagnostic criteria, 60% to more than 90% of dementias are attributed to Alzheimer's disease. That breadth is itself informative: case mix changes across community, clinic and institutional settings, while diagnostic boundaries alter the numerator. Alzheimer disease may be the leading contributor without being the only cause in an individual. Mixed pathology is common clinically, but its detailed diagnosis and management belong elsewhere in the chapter.

4.3 Pathology without dementia: the denominator problem

A crucial conceptual distinction is that biological change may precede, or never produce, the clinical syndrome. As many as 70% of cognitively normal older adults show clear Alzheimer histopathological changes, and up to 30% meet full autopsy-confirmed neuropathological criteria. These observations do not make clinical assessment obsolete. They show that pathological burden, cognitive reserve, co-pathology and resilience influence whether tissue-level disease becomes manifest as functional cognitive decline.

RULE **Separate clinical prevalence from pathological prevalence.** Alzheimer-type change in a cognitively normal person is not equivalent to dementia, while dementia remains a clinical syndrome whose cause requires appropriate diagnostic evaluation.

This distinction also protects interpretation of biomarkers and genetic results. A marker that shifts probability cannot be read in isolation from cognition, function and longitudinal course. Conversely, normal performance at one encounter does not settle the lifetime biological question. The examiner is looking for this layered model: pathology can exist without current clinical impairment, risk can exist without pathology being demonstrated, and a clinical syndrome can have more than one pathological contributor.

The numerical contrast is best used to explain penetrance and expression rather than to claim that silent pathology is harmless. Some individuals may be earlier in a disease process; others may tolerate considerable burden without crossing a clinical threshold. The neuropathological lesions and biomarker interpretation are owned by their dedicated sections. Here the epidemiological lesson is that the denominator must match the question being asked: clinical cases, pathological cases and at-risk persons are different populations.

4.4 Modifiable and non-modifiable risk

Age is the strongest identified risk factor. Genetic predisposition is also substantial, with 60% to 80% of attributable risk assigned to genetic influence in the cited synthesis. Neither statement means that environment is negligible. A high heritable contribution describes variation in a population under its prevailing conditions; it does not tell an individual that outcome is fixed, and it does not make preventable exposures clinically irrelevant.

The 2020 report of the Lancet Commission estimated that **about 40% of dementias worldwide may be related to modifiable factors**. This estimate concerns dementia broadly, not Alzheimer disease alone, so its scope must be retained. The listed factors span the life course: less education; midlife hypertension and high cholesterol; cardiovascular or cerebrovascular disease; diabetes and obesity; depression across earlier and later life; sleep disturbance; traumatic brain injury; excessive alcohol use; smoking; anticholinergic medication burden; physical inactivity; air pollution; social isolation; and hearing impairment.

DSM-5-TR identifies low educational status, midlife hypertension, obesity, hearing loss, late-life smoking, depression, physical inactivity, social isolation and diabetes as Alzheimer-specific risk and prognostic factors. It also states that **clustering of vascular risk factors** increases risk. Traumatic brain injury, particularly in men, may increase risk, although the association remains controversial. This wording matters: a possible association should not be upgraded into settled causation.

- **Life-course accumulation:** education, vascular exposure, sensory loss and social context act over different periods rather than at one late-life moment.
- **Shared pathways:** several factors affect both vascular and neurodegenerative outcomes, so risk reduction need not depend on assigning a single future dementia subtype.

- **Clustering:** multiple vascular risks together carry more concern than an isolated label viewed without context.
- **Association is not destiny:** a risk factor changes probability; it does not establish present disease or guarantee future disease.
- **Scope discipline:** the broad dementia-attributable estimate and the Alzheimer-specific factor list should be presented as related but not identical constructs.

GOOD TO REMEMBER

Risk review is clinically useful even when causation cannot be assigned in one patient. Address vascular health, hearing, activity, mood, social connection, alcohol, smoking and medication burden because they affect present health and may also alter future cognitive risk. Do not promise that modifying one factor will prevent Alzheimer disease in that individual.

4.5 Genetic architecture: cause, susceptibility and background

The genetics become clear when divided into categories. **Familial Alzheimer disease** is described by Kaplan and Saddock as accounting for about **5%** of cases, with early onset, autosomal dominant inheritance and high penetrance. The remaining sporadic forms have a genetic contribution resembling that of other common neuropsychiatric disorders: many susceptibility variants interact with age and non-genetic exposures rather than acting as a single sufficient cause.

Twin evidence supports a meaningful but incomplete inherited contribution. Concordance in **monozygotic twin pairs** is about 50%. If genes were irrelevant, such concordance would be difficult to explain; if genes alone were sufficient for all disease, concordance would be complete. The gap between those possibilities points toward variable penetrance, polygenic background, environmental exposure, survival and stochastic biological processes.

MNEMONIC

Cause, Susceptibility, Context

Cause: rare dominant mutations with high penetrance. **S**usceptibility: common alleles that alter probability. **C**ontext: age, vascular and life-course exposures that influence whether risk becomes clinical disease.

TRAP **Do not quote one universal percentage for "familial Alzheimer disease".** The supplied texts use different boundaries: one gives about 5% for familial forms, while another describes clearly inherited early-onset disease as a smaller minority. State the source and definition, and do not pretend the discrepancy has been resolved.

The conflict is not merely editorial. "Familial" may mean any aggregation among relatives, or it may be reserved for a pedigree showing high-penetrance autosomal dominant transmission. "Early onset" is also not a synonym for "Mendelian": causative mutations are enriched in earlier presentations but are not present in every such patient. In an answer, define the category before

quoting its proportion and keep mutation-caused disease separate from common familial clustering.

4.6 The autosomal dominant genes: APP and the presenilins

The first identified cause of familial early-onset disease involved the **amyloid precursor protein gene (APP)** on chromosome 21. Pathogenic variants increase extracellular longer beta-amyloid fragments, including amyloid-beta 42. The high-yield link is therefore direct: the encoded precursor sits upstream of amyloid generation, and a mutation can shift processing toward the aggregation-prone product. The fuller amyloid cascade and neuropathological consequences belong in the dedicated pathology section.

Presenilin-1 and presenilin-2 are located on chromosome 14q24.3 and chromosome 1q, respectively. They are integral transmembrane proteins with at least seven transmembrane domains and function as aspartyl proteases in gamma-secretase cleavage of the precursor protein, generating beta-amyloid. Thus, these dominant genes converge on production or processing of the same molecular substrate even though their chromosomal locations differ.

GENETIC CATEGORY	LOCATION OR SCALE	BIOLOGICAL OR CLINICAL MEANING
APP	Chromosome 21	Dominant familial cause; pathogenic variants increase longer extracellular beta-amyloid fragments
PSEN1	Chromosome 14q24.3	Presenilin component of gamma-secretase processing
PSEN2	Chromosome 1q	Presenilin component of gamma-secretase processing
APOE	Chromosome 19q	Common late-onset susceptibility locus, not a deterministic dominant mutation
Genome-wide association loci	More than 40 confirmed risk genes	Common small-effect architecture converging on several biological pathways

Kaplan and Saddock's edition lists 185 PSEN1 mutations, 13 PSEN2 mutations and 24 APP mutations, virtually all transmitted as autosomal dominant, fully penetrant variants. These are edition-specific counts rather than timeless biological constants; mutation catalogues change as new variants are characterised. For examination purposes, the durable distinction matters more than memorising a potentially moving inventory: these genes can cause high-penetrance familial disease, whereas common susceptibility alleles modify risk.

4.7 APOE: the major susceptibility locus

The **apolipoprotein E gene (APOE)** lies on chromosome 19q. Its three recognised alleles are epsilon 2, epsilon 3 and epsilon 4, distinguished by amino acids at residues 112 and 158: epsilon 2 is Cys112/Cys158, epsilon 3 is Cys112/Arg158, and epsilon 4 is Arg112/Arg158. Epsilon 3 is

the most common allele in most populations. The structural detail is worth knowing, but its clinical meaning is the real examination target.

APOE epsilon 4 is enriched in late-onset disease. Its frequency is approximately 50% in familial late-onset Alzheimer disease and 40% in sporadic late-onset disease, compared with about 16% in controls. Across late-onset cases, 30% to 60% carry at least one epsilon 4 allele. Because many affected people do not carry it, and many carriers do not inevitably become affected, it is a susceptibility allele rather than a diagnostic marker or a Mendelian cause.

The dose-risk relation is strong but still probabilistic. **One copy raises risk about 4-fold**, while two copies raise it about 10-fold. Relative risk is not the same as absolute lifetime probability. The latter depends on background population risk, age, ancestry, competing mortality and other genetic and environmental influences. A result therefore cannot be translated into a certain personal forecast by repeating the fold-change alone.

> **PITFALL**
> Do not tell a patient that an APOE epsilon 4 result diagnoses Alzheimer disease or guarantees future dementia. It changes susceptibility. Interpretation requires the clinical phenotype, family history, purpose of testing and appropriate genetic counselling; the relative-risk figure alone is not an individual prognosis.

4.8 Beyond APOE: modifiers, pathways and genetic convergence

GAB2 illustrates gene-gene interaction. It was identified through genome-wide association as an additional risk allele among APOE epsilon 4 carriers, but not among affected non-carriers. Carrying both risk alleles was associated with an almost 25-fold greater risk than carrying neither. Mechanistically, GAB2 can modulate beta-amyloid production through interaction with GRB2, which physically interacts with APP and the presenilins. This is a modifier model, not a replacement for the distinction between causal mutations and susceptibility alleles.

The wider **polygenic risk architecture** includes more than 40 confirmed risk genes identified through genome-wide association studies. Unlike the dominant genes, each generally has a small effect, with an odds ratio below 1.2, and is relatively common in the population. Their biological convergence is more informative than any isolated locus: implicated pathways include immunity, lipid metabolism, tau binding and amyloid precursor protein metabolism. Common disease emerges from networks rather than one obligatory mutation.

The association between Down syndrome and Alzheimer disease is an important genetic bridge, but its gene-dosage mechanism, diagnostic classification and assessment of cognitive decline are taught in the dedicated Down-syndrome section. Intellectual disability as a developmental syndrome belongs in the Intellectual disability and autism spectrum disorder notes. This boundary matters because a shared molecular link does not make the baseline developmental condition and later cognitive decline interchangeable clinical constructs.

The complete genetic answer therefore ends with hierarchy. Rare dominant variants in the APP gene and presenilin genes can be highly penetrant causes. The APOE epsilon 4 allele is a major

common susceptibility factor with a copy-related increase in risk. Other common loci contribute smaller effects and reveal convergent biological pathways. Age and modifiable exposures remain clinically relevant throughout. None of these layers should be allowed to masquerade as another.

> Rare APP and presenilin mutations can cause autosomal dominant Alzheimer disease; APOE epsilon 4 modifies late-onset risk but is neither necessary nor sufficient for the disorder.

5 · Alzheimer's disease: neuropathology and the amyloid cascade

The scoring answer links lesion, location, sequence and mechanism. Alzheimer pathology is not merely a list of plaques, tangles and transmitter deficits. The lesions have different cellular compositions, begin in different regions and acquire different explanatory roles. A mature answer therefore separates the amyloid map from the tau map, relates both to neuronal and synaptic injury, and then treats the amyloid cascade as a mechanistic model rather than as a synonym for the microscopic diagnosis.

The practical value of this distinction is interpretive. A pathologic finding matters because of what it contains, where it occurs, how widely it has spread and what other processes accompany it. The same discipline prevents two common errors: assuming that every lesion has the same spatial progression, and treating one proposed causal chain as if it made the rest of the pathology redundant.

LESION	MAP	SYSTEM	MODEL
IDENTIFY THE STRUCTURE	FOLLOW ITS SPREAD	TRACE CIRCUIT LOSS	TEST THE CAUSAL CLAIM

5.1 The two signature lesions are not the same map

Amyloid plaques are the defining lesion in the standard neuropathologic account. They consist of insoluble deposits of fibrillar amyloid-beta, with dystrophic neurites, microglial cells and reactive astrocytes arranged around the deposit. **Neurofibrillary tangles**, by contrast, are intracellular inclusions. Their paired helical filaments contain aggregated, hyperphosphorylated microtubule-associated tau protein. Composition already gives the first conceptual separation: one lesion is organised around fibrillar amyloid-beta deposition, while the other reflects abnormal tau aggregation within neurons.

The second separation is topographic. Plaques first appear widely through the neocortex and are found later in the entorhinal cortex and hippocampus, followed by subcortical regions. Tangles begin in the entorhinal cortex, extend to the hippocampus and lateral temporal lobe, and only then become widespread through the neocortex. Thus, “early Alzheimer pathology” is incomplete unless the candidate states which lesion is being mapped.

AXIS	AMYLOID PLAQUE	NEUROFIBRILLARY TANGLE
Basic structure	Insoluble fibrillar amyloid-beta deposit with surrounding dystrophic neurites and glial reaction	Intracellular inclusion containing paired helical filaments of aggregated hyperphosphorylated tau
Earliest distribution	Neocortex	Entorhinal cortex
Subsequent spread	Entorhinal cortex and hippocampus, then subcortical regions	Hippocampus and lateral temporal lobe, then wider neocortex
Diagnostic caution	Defining lesion in this pathologic account	Not pathognomonic, because tangles are also the principal lesion in MAPT-related frontotemporal dementias

MNEMONIC

A-N, T-E

Amyloid begins in the **N**eocortex; **T**au begins in the **E**ntorhinal cortex. The letters prevent the classic reversal, while the full spread sequence still needs to be understood.

RULE **Always name the lesion before describing its spread.** Neocortex-first describes the plaque map; entorhinal-first describes the tangle map.

5.2 Gross progression and the regional burden of neuronal loss

Macroscopic progression does not become conspicuous at the same moment as microscopic disease. In early Alzheimer disease, the brain may look grossly normal, or atrophy may be confined to the hippocampus. One macroscopic clue to hippocampal volume loss is enlargement of the inferior horn of the lateral ventricle. This observation matters because a subtle gross specimen does not negate an active microscopic process.

With progression, the macroscopic pattern becomes more evident: cortical sulci widen and the ventricles enlarge. Atrophy is often particularly prominent in posterior temporal, parietal and frontal lobes. These findings express tissue loss at the organ level; they should not be used to collapse the distinct plaque and tangle sequences into a single map.

At the cellular level, **neuronal loss** is most marked in the entorhinal cortex and hippocampus. This regional emphasis helps explain why the medial temporal system is central to the pathologic narrative even though amyloid plaques are described as appearing first across neocortex. The lesion that appears first, the region with greatest neuronal loss and the structure most closely aligned with cognitive evolution need not be identical constructs.

Later disease also involves neurotransmitter systems. Cholinergic neurons in the nucleus basalis are lost, and brainstem noradrenergic and serotonergic neurons are affected. The pathology therefore expands from local tissue injury to disconnection of projection systems. This is the

bridge from morphologic description to neurochemical models: circuit failure is real, but it does not by itself identify the initiating event.

GOOD TO REMEMBER

When shown a gross specimen, describe what is visible before inferring cause: hippocampal atrophy and inferior-horn enlargement may be early clues; later, widened sulci, ventricular enlargement and distributed lobar atrophy become clearer.

5.3 Tau staging: an ordered account of neurofibrillary spread

Braak and Braak staging organises neurofibrillary change by its anatomic reach. The sequence is **stages I-II, transentorhinal; stages III-IV, limbic; and stages V-VI, isocortical or neocortical**. It is a staging system for neurofibrillary pathology, not a relabelling of the plaque sequence. Its educational value lies in converting a static slide into an evolving regional process and thereby supporting an anatomic understanding of changing cognitive deficits.

Photomicrographic comparisons make the burden dimension concrete. In **stage VI, plaques and tangles are extensive throughout the neuropil, whereas stage IV shows some tangles with substantially less neurofibrillary pathology in the neuropil**. The comparison preserves direction: the more advanced stage carries the heavier and more widespread pathologic burden. Isolated tangles in the entorhinal cortex without plaques can also be seen in normal ageing. Their presence alone cannot therefore be read as a complete Alzheimer pattern.

Two interpretive boundaries follow. First, staging is about distribution and burden, not merely whether a tangle can be found. Second, a tangle is not uniquely diagnostic of Alzheimer disease. MAPT-related frontotemporal dementias also have tangles as their principal lesion. The diagnosis therefore depends on the configuration of lesions and regions, not on a decontextualised microscopic label.

TRAP **Do not write that plaques and tangles both begin in the hippocampus.** The hippocampal and entorhinal emphasis in neuronal loss can seduce candidates into assigning the same starting point to both lesions.

5.4 Additional lesions and co-pathology

Real brains need not display a chemically pure single-protein disorder. **Alpha-synuclein co-pathology** may be detected in **up to 50%** of individuals with Alzheimer disease. It most commonly appears as thread-like deposits or intraneuronal Lewy bodies in the amygdala and transentorhinal cortex, although Lewy bodies may also occur in neocortex. The correct lesson is heterogeneity of pathologic burden, not automatic reassignment of the clinical diagnosis.

Other microscopic findings include **granulovacuolar degeneration and Hirano bodies**. They most commonly affect pyramidal neurons in the hippocampus and may occur in Alzheimer disease. These findings enrich the neuropathologic description but should remain subordinate to the central organisation of plaques, tangles, regional neuronal loss and projection-system involvement.

This layered view also improves clinicopathologic reasoning. A patient may have a dominant Alzheimer pattern with additional protein aggregation. Conversely, finding one ancillary lesion does not tell the examiner how much amyloid deposition, tau spread or neuronal loss is present. The disciplined report separates dominant pattern, stage and co-pathology instead of forcing every observation into an all-or-none label.

PITFALL

A pathology report that mentions Lewy bodies alongside Alzheimer lesions is not self-interpreting. Do not ignore the co-pathology, but do not replace the dominant diagnosis solely because an additional inclusion is present; distribution and overall burden remain essential.

5.5 The cholinergic account: strong neurochemical basis, limited causal reach

A transmitter deficit can explain impairment without explaining the whole disease. The **cholinergic hypothesis** rests on converging neuropathologic and neurochemical observations: a major decline in basal-forebrain cholinergic projections to cortex, reduced choline acetyltransferase activity and marked loss of cholinergic cell bodies in the nucleus basalis. Reductions in enzyme activity and nucleus-basalis integrity correlate with cortical neuritic plaque density, and cholinergic deficits correlate with poorer cognitive-test performance.

The model therefore has genuine explanatory force. It connects a defined projection system, a measurable neurochemical deficit, cortical pathology and cognitive performance. It also supplies a rational bridge between tissue pathology and symptomatic impairment. Yet correlation between circuit loss and cognitive deterioration does not, by itself, establish that cholinergic depletion initiated the entire disease process.

The later formulation is appropriately narrower: progressive loss of cholinergic neurons and falling acetylcholine levels contribute to cognitive deterioration, but cholinergic depletion alone is no longer regarded as sufficient to account for all Alzheimer symptoms. Nor is the cholinergic system the only network affected. Noradrenergic projections from the locus ceruleus to cortex and glutamate systems are also involved, alongside the serotonergic and other brainstem losses described in the neuropathology.

This is why the cholinergic account should not be set against every molecular process as an exclusive rival. It addresses an important level of the disease, namely transmitter and projection-system failure. Amyloid production and deposition, tau phosphorylation, inflammatory responses and neuroprotective processes address more fundamental questions about how neurodegeneration develops and propagates.

5.6 The amyloid cascade: from precursor processing to tissue injury

The amyloid cascade hypothesis proposes a causal sequence rather than simply naming a lesion. Faulty processing of amyloid precursor protein generates neurotoxic amyloid-beta fragments. These fragments are then placed upstream of neurodegeneration and other hallmarks of Alzheimer pathology, including phosphorylated tau aggregation and synaptic loss. The

strength of the model is its attempt to connect molecular processing, deposited protein, cellular reaction and downstream structural injury in one chain.

In the classical processing account, amyloid precursor protein is cleaved by alpha-, beta- and gamma-secretases. Alpha-secretase cleavage yields a nontoxic metabolite. The alternative sequence, beta-secretase followed by gamma-secretase, produces **amyloid-beta 42**, described as relatively insoluble. It fibrillises into diffuse plaques, provokes an inflammatory response and is neurotoxic. Proposed downstream effects include axonal breakdown, the formation of neuritic plaques and neurofibrillary-tangle formation in surviving neurons.

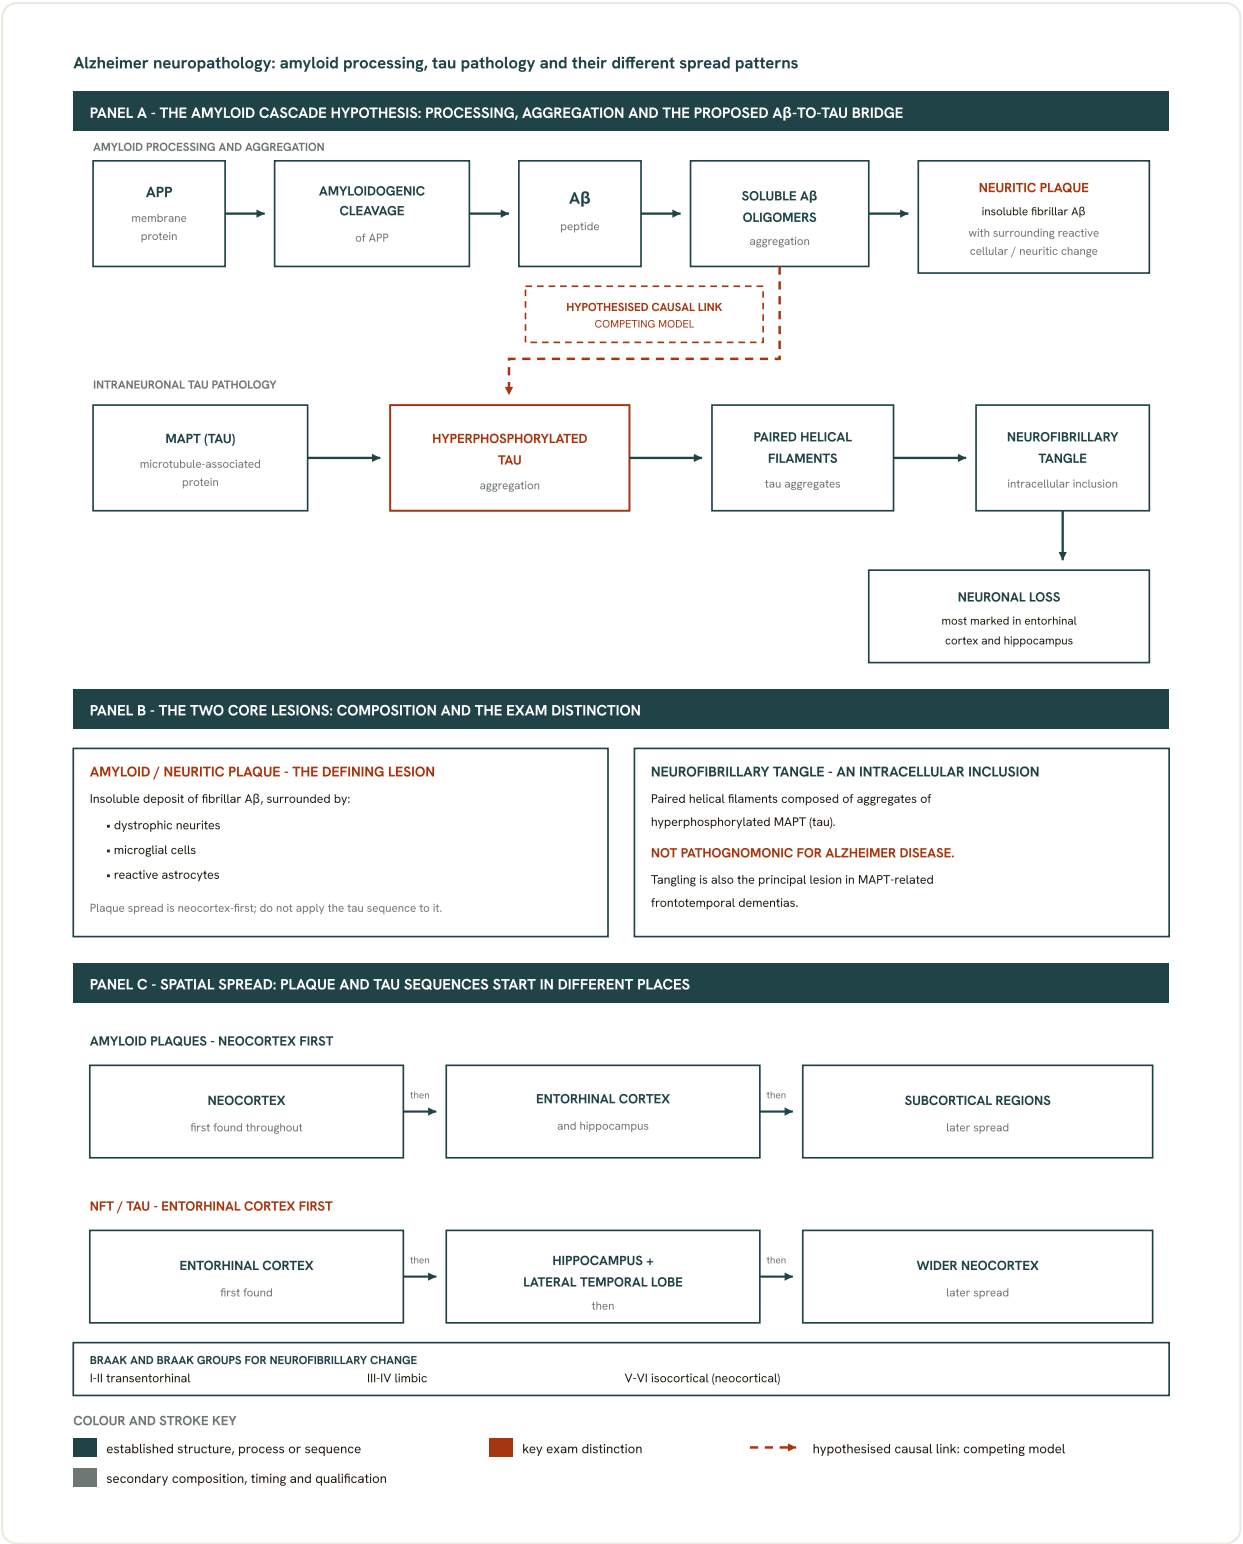


Fig 32.2 **Alzheimer neuropathology and the amyloid cascade hypothesis.** APP processing and Aβ aggregation are shown above tau hyperphosphorylation, paired helical filaments and neurofibrillary tangles. The lower panel keeps the neocortex-first amyloid plaque sequence separate from the entorhinal-first tau sequence. (Kaplan & Saddock's CTP 2024; Hodges, Cognitive Assessment for Clinicians.)

The accompanying figure should be read as a directional hypothesis. It begins with processing of the precursor protein, separates the nontoxic route from the amyloid-producing route, and follows the latter toward aggregation, inflammation and neuronal injury. Its arrows express proposed causal ordering. They do not mean that every downstream process is passive, that all

lesions share the same anatomic spread, or that pathology can be inferred from a single node in isolation.

The word “cascade” is useful precisely because it invites questions at each link. Is the relevant fragment generated? Does it remain soluble or aggregate? Does aggregation recruit glial and inflammatory responses? How are axonal and synaptic injury related to tau change? A high-quality account states the sequence, then examines the evidence and the limitations rather than repeating the pathway as settled dogma.

5.7 Why amyloid was placed upstream

Several observations are used to support an upstream role for amyloid precursor processing. They are best treated as converging lines of evidence, not as interchangeable proofs. Their relevance here is mechanistic; the broader epidemiology, penetrance and counselling implications belong to the dedicated section on Alzheimer epidemiology, risk factors and genetics.

- People with Down syndrome who survive past 40 years almost always develop Alzheimer disease in the cited account, linked to an extra copy of the APP gene.
- Rare inherited Alzheimer disease can result from mutations in APP, directly connecting altered precursor biology with disease.
- Presenilin interacts with gamma-secretase, while rare inherited disease is associated with presenilin 1 or presenilin 2 mutations.
- The epsilon-4 form of apolipoprotein E increases the rate of beta-amyloid production in the cited mechanistic account.

Taken together, these observations connect gene dosage, rare pathogenic variants, secretase-associated biology and a susceptibility-associated apolipoprotein form with amyloid production or processing. The inference is that amyloid biology is not merely an incidental end product of dying tissue. However, support for an upstream role does not establish that amyloid alone determines the pace, topography or clinical expression of late-onset disease.

The causal claim should therefore be stated with correct scope. Genetic observations strongly motivate the pathway, especially where precursor processing is directly altered. They do not license the candidate to substitute a genetics essay for neuropathology, nor to assume that every sporadic case follows a simple single-gene logic. Mechanistic relevance and population-level determinism are different questions.

5.8 Tau, inflammation and the limits of a pure cascade

The major challenge is not whether amyloid belongs in Alzheimer pathology; it plainly does. The challenge is whether a simple amyloid-first chain is sufficient to explain the disease. Clinical trials of immunotherapies directed at amyloid-beta 1-42 have failed in the evidence base summarised by the source. Possible explanations include intervention after the disease process had become too advanced and enrolment of participants who did not actually have amyloid pathology. These possibilities preserve the hypothesis from a premature dismissal, but they also expose how strongly trial interpretation depends on timing and biologic selection.

Late-onset Alzheimer disease is consequently better approached as complex. The molecular links within the cascade require fuller explanation, and associated processes, especially inflammation, need attention. The glial reaction around plaques and the inflammatory response provoked in the processing model are not decorative histology. They indicate that deposited protein interacts with surrounding tissue and that injury may be mediated through more than direct peptide toxicity.

Tau-mediated injury also has active mechanistic significance. Amyloid-beta contributes to hyperphosphorylation of microtubule-associated tau, promoting its aggregation into neurofibrillary tangles. Once established, tau pathology causes synapse loss and neuronal death independently of amyloid-beta effects. Tau is therefore not merely a passive stain recording prior amyloid injury; it can carry damaging consequences of its own.

The resulting model is neither a rejection of the cascade nor an excuse to label every process an independent hypothesis. Amyloid production and deposition, tau phosphorylation and aggregation, inflammation, synaptic loss, neuronal death and transmitter-system failure occupy related but distinguishable levels. A useful synthesis asks what each level explains and where causal certainty weakens.

EXPLANATORY LEVEL	WHAT THE EVIDENCE SUPPORTS	WHAT MUST NOT BE OVERCLAIMED
Cholinergic deficit	Projection loss, reduced choline acetyltransferase activity and association with cognitive performance	Cholinergic depletion alone does not account for the whole syndrome
Amyloid cascade	Faulty precursor processing is placed upstream of toxic fragments, neurodegeneration, phosphorylated tau aggregation and synaptic loss	A simple late-onset model is challenged by treatment failures, timing problems and biologic complexity
Tau process	Amyloid-beta promotes tau hyperphosphorylation and aggregation; tau can then drive synaptic and neuronal injury independently	Tau should not be reduced to a passive downstream marker
Multiple-system degeneration	Amyloid deposition, tau phosphorylation, inflammation and several transmitter systems are implicated	No single transmitter deficit should be mistaken for the complete pathogenesis

5.9 An integrated answer for pathology and pathogenesis

The final synthesis should move from observation to inference in a fixed order. Begin with the gross brain: it may be normal-looking early, may show hippocampal atrophy with inferior-horn enlargement, and later shows wider cortical and ventricular change. Move next to the core microscopic lesions, keeping plaque composition and neocortex-first spread separate from tangle

composition and entorhinal-first spread. Then add regional neuronal loss, neurotransmitter-system degeneration and the possibility of ancillary lesions or co-pathology.

Only after that description should pathogenesis be introduced. The cholinergic account explains a clinically important projection-system deficit but not the entire disease. The amyloid cascade places abnormal precursor processing and neurotoxic amyloid fragments upstream. Tau aggregation, inflammatory responses and synaptic injury then complicate any purely linear reading, especially because tau can cause further loss independently of amyloid-beta.

This structure also keeps adjacent chapter material in its proper place. Neuropathology does not require a detour into diagnostic criteria, biomarker cut-offs, imaging physics, current antibody regimens or symptomatic drug doses. Those subjects use the pathology developed here, but they answer different questions. The pathology section asks what is in the brain, where it spreads and how the proposed molecular sequence relates to the observed injury.

A concise conclusion can therefore be both balanced and decisive: Alzheimer disease has characteristic amyloid and tau pathology with different spatial trajectories, accompanied by selective neuronal and projection-system loss. The amyloid cascade offers a powerful upstream framework, supported by processing and genetic observations, but the mature model must accommodate active tau toxicity, inflammation, co-pathology and multisystem neurodegeneration.

> Memorise the separation: amyloid plaques begin in neocortex, neurofibrillary tangles begin in entorhinal cortex, and the amyloid cascade is a causal hypothesis that must be integrated with active tau injury, inflammation and transmitter-system loss.

6 · Alzheimer's disease: clinical staging, diagnostic criteria and management

Why it matters clinically. Alzheimer's disease is not recognised by a single forgotten appointment, a single low screening score, or an image interpreted in isolation. The examiner expects a coherent account that moves from the characteristic clinical trajectory to the neurocognitive syndrome, establishes the likely aetiology, grades diagnostic certainty, and then links the stage to a proportionate care plan. It separates preserved ability from expected progression.

The central idea is progression with changing cognitive and functional emphasis. Early disease may be dominated by an amnestic syndrome while ordinary self-care remains intact. Dementia becomes the appropriate syndrome label when cognitive decline interferes with independent functioning. Later, multiple cognitive domains, behaviour, mobility, continence, communication, and physical dependence become involved. Atypical presentations complicate recognition but do not abolish the need to demonstrate a compatible progressive course and exclude a better explanation.

Memory	Function	Course	Care
EARLY SIGNAL	SYNDROME BOUNDARY	DIAGNOSTIC WEIGHT	STAGE MATCHED

6.1 The clinical trajectory and the meaning of stage

A useful bedside model describes a predictable progression through **three stages**: mild cognitive impairment, mild-to-moderate dementia, and advanced dementia. The model is clinical rather than merely chronological. It asks which cognitive systems are failing, whether everyday independence is preserved, which neuropsychiatric and physical problems have appeared, and how much supervision is now required. It also corrects a common misconception: typical Alzheimer's disease does not begin as simultaneous global intellectual failure.

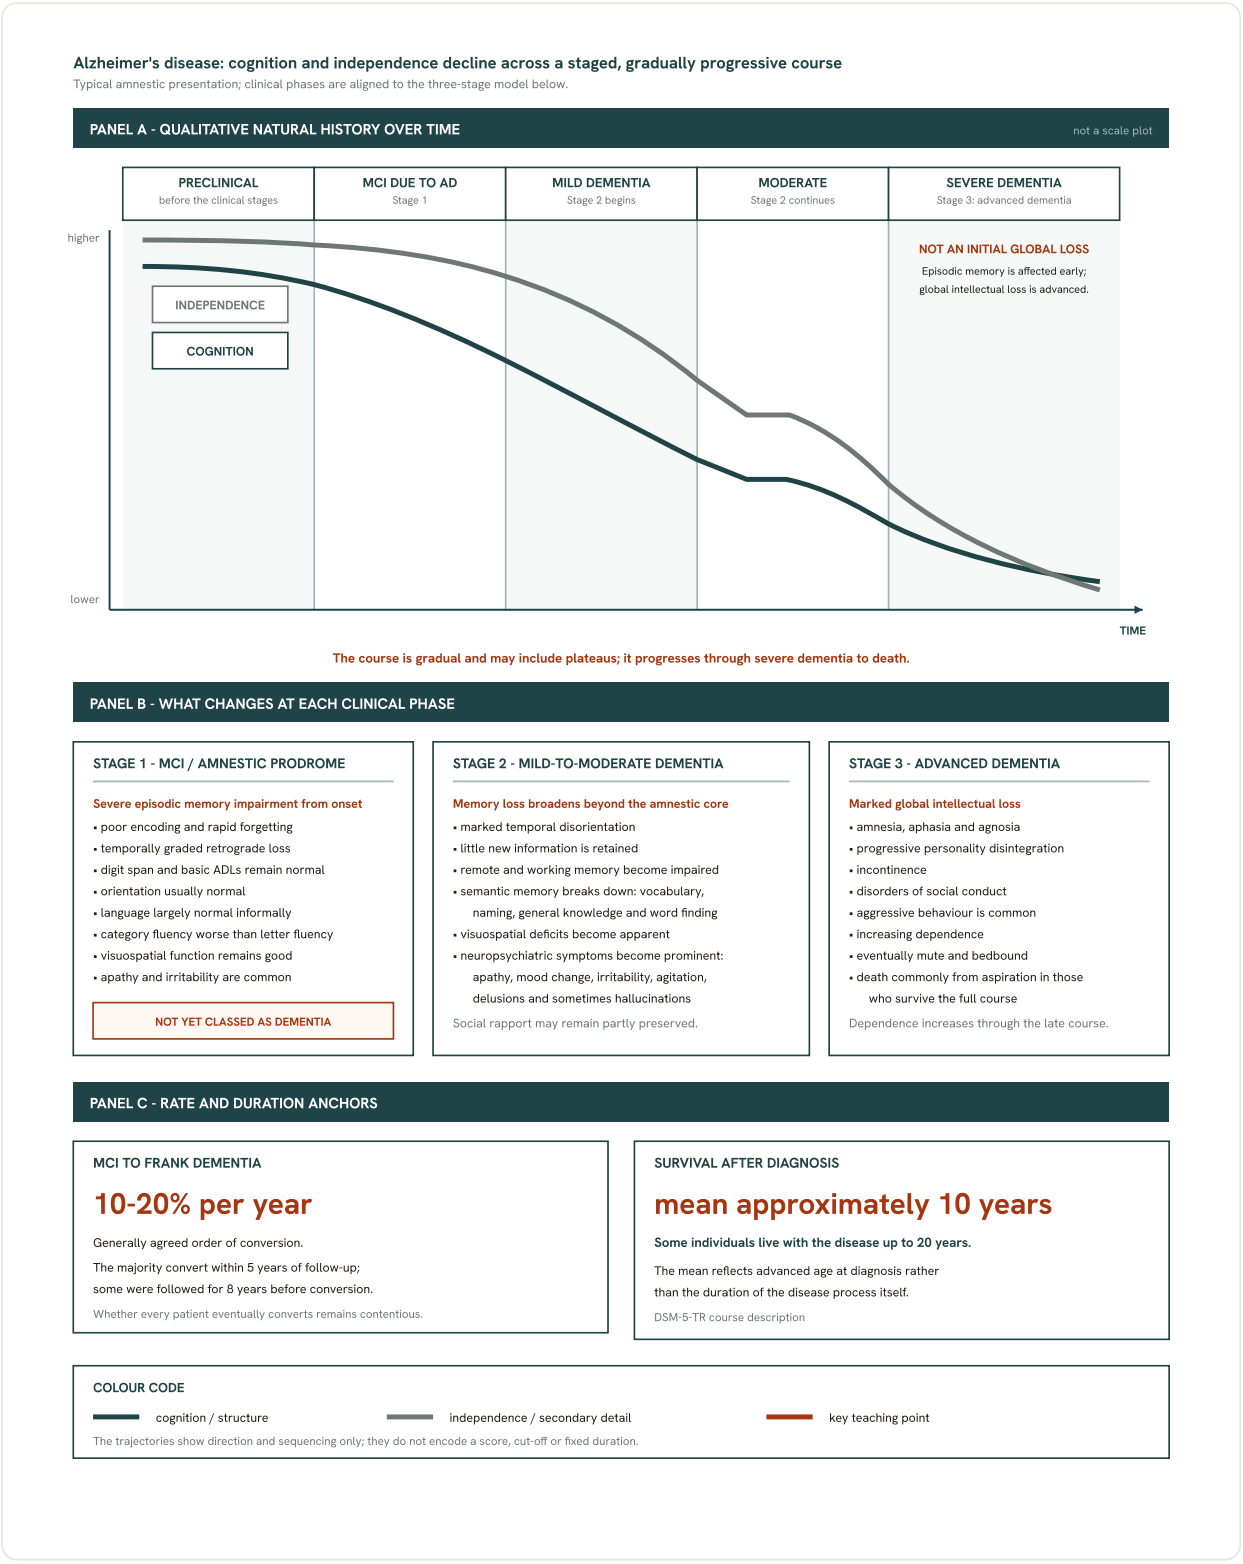


Fig 32.3 Natural history of typical Alzheimer's disease. Cognitive decline begins before major functional dependence, progresses from the amnestic MCI prodrome through mild-to-moderate dementia, and culminates in advanced global impairment and dependence.

The accompanying figure should be read as a continuum rather than a set of sealed compartments. Cognitive decline and functional loss do not move in perfect lockstep. A highly supported patient may appear functionally better than the underlying impairment suggests, while sensory loss, low education, physical illness, depression, or an unfamiliar testing context can make performance look worse. Staging therefore needs history from the patient and a

reliable informant, direct cognitive examination, and a functional account grounded in the person's usual roles.

RULE **Stage the person by cognition plus function, not by a score alone.** The clinically decisive question is whether decline has crossed from preserved independence into dependence, and which new risks that transition creates.

Preclinical biological change, symptomatic mild cognitive impairment, and dementia can be represented along a unified natural-history line, but they are not interchangeable diagnostic labels. Biological evidence can increase confidence about aetiology; it does not replace the clinical description of the person's current syndrome. Conversely, an apparently normal brief screen does not neutralise a convincing history of selective episodic-memory failure in the prodromal phase.

6.2 The staged clinical phenotype

CLINICAL STAGE	COGNITIVE PROFILE	FUNCTION AND BEHAVIOUR	INTERPRETIVE POINT
Stage 1: mild cognitive impairment	Severe episodic-memory impairment is present from the outset, with poor encoding, rapid forgetting, anterograde amnesia, and a temporally graded retrograde loss. Orientation is usually preserved. Informal language and visuospatial performance may remain largely satisfactory, although category fluency may be weaker than letter fluency.	Basic activities of daily living remain normal. Apathy and irritability may appear, but loss of functional independence is not yet sufficient for dementia.	A normal brief cognitive screen does not exclude this selective amnestic prodrome. General administration and scoring of cognitive screens belong in the Psychometry, Assessment and Descriptive Psychopathology notes.
Stage 2: mild-to-moderate dementia	Worsening memory and attention produce marked temporal disorientation and poor retention of new information. Remote and working memory become impaired. Semantic breakdown produces reduced vocabulary, word-finding difficulty, semantic paraphasias, impaired naming, and loss of general knowledge. Visuospatial deficits become evident.	Neuropsychiatric symptoms become more prominent, including apathy, mood disturbance, irritability, agitation, delusions, and sometimes hallucinations. Social rapport may remain partly preserved.	The syndrome is no longer a circumscribed memory complaint. Cognitive decline now compromises everyday functioning and spreads across domains.
Stage 3: advanced dementia	There is marked global loss, with amnesia, aphasia, and agnosia. Communication may ultimately be lost.	Personality disintegrates progressively. Incontinence, disordered social conduct, aggressive behaviour, immobility, and increasing dependence become prominent.	Management priorities shift toward comfort, prevention of avoidable suffering, carer support, and anticipatory planning.

The early memory disorder is not merely inefficient retrieval. Poor encoding and rapid forgetting mean that cues may provide limited rescue, and the patient cannot reliably retain the substance of a conversation. This explains repeated questions, missed arrangements, and loss of newly learned routes even when social manner, vocabulary, and simple attention look deceptively intact. As disease advances, failure of semantic knowledge, visuospatial processing, and executive control progressively changes both test performance and real-world competence.

GOOD TO REMEMBER

Ask for an example of lost independence, not only a list of cognitive symptoms. Errors in medication use, money management, navigation, cooking, communication, or occupational tasks often define the syndrome transition more clearly than a total screening score.

6.3 Atypical Alzheimer's disease and diagnostic flexibility

The amnestic presentation is the majority pattern, but Alzheimer's pathology can first become clinically visible through another network. The principal atypical variants include language-led and posterior cortical presentations. **Logopenic progressive aphasia** presents with word-finding hesitancy and pauses, phonological errors, and reduced digit or sentence span. Knowledge of word meaning is relatively intact, while apraxia of speech and agrammatism are absent. The language phenotype therefore reflects impaired access and phonological working memory rather than early loss of semantic knowledge or motor-speech programming.

Posterior cortical atrophy begins with higher visual and visuospatial difficulty. The patient may become lost in unfamiliar surroundings or struggle to reach and grasp accurately. Posterior cortical atrophy may be visible on structural imaging, while memory, language, and insight remain relatively preserved early. Severe progression may produce functional blindness. A rarer frontal-predominant presentation can feature apathy without the severe early amnesia expected in the typical form.

The practical lesson is not to force every patient into an amnestic template. A progressive focal cortical syndrome may still be due to Alzheimer's disease, but its first deficit should be documented precisely and competing neurodegenerative, vascular, structural, psychiatric, sensory, and systemic explanations considered. Imaging may support the anatomical pattern, while the physics and general interpretation of imaging modalities are covered in the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

TRAP **"Memory is not the presenting complaint, so this cannot be Alzheimer's disease."** That statement misses language-led and posterior cortical presentations. State the dominant network syndrome first, then judge whether its course and aetiological evidence support Alzheimer's disease.

6.4 Conversion, progression and prognosis

For amnestic mild cognitive impairment, a generally accepted conversion estimate is **10-20% per year**. The majority who convert do so during longitudinal follow-up, although later conversion occurs and it remains contentious whether every patient will eventually develop dementia. This uncertainty matters when counselling. Mild cognitive impairment is a risk state and a clinical syndrome, not a promise of inevitable dementia in a particular individual.

Published conversion figures vary because clinical referral samples, community cohorts, definitions, ascertainment, and duration of follow-up differ. The correct interpretation is therefore not to transplant a study-specific proportion into an individual prognosis. Instead, communicate that risk is meaningfully increased, arrange longitudinal review, and reassess both

cognition and function. Change within the same person is usually more informative than a single isolated performance.

Once major or mild neurocognitive disorder due to Alzheimer's disease is established, the course is generally gradual and progressive, although plateaus may occur. Mean survival after diagnosis is approximately **10 years**, and some people live much longer. The estimate is shaped strongly by the advanced age at which diagnosis is often made and should not be presented as a countdown. In the late phase, patients may become mute and bedbound; among those who survive the full course, aspiration is the commonest terminal event.

Prognostic conversation should therefore be staged and revisited. Early discussion concerns uncertainty, preserved abilities, adaptation, safety, and future preferences. With increasing dependence, the focus expands to supervision, caregiver capacity, behavioural symptoms, swallowing and mobility problems, burdens of intervention, and goals of care. The point is preparedness without therapeutic nihilism: progression is expected, but suffering, unsafe care, and unsupported carers are not inevitable consequences that must simply be accepted.

6.5 Diagnostic architecture: syndrome, course, cause and exclusion

Diagnosis works best as a layered argument. First establish a major or mild neurocognitive syndrome and describe the domains affected. Next show an insidious onset and gradual progression. Then decide how strongly the clinical and available biological evidence attributes the syndrome to Alzheimer's disease. Finally, test whether vascular disease, another neurodegenerative disorder, substance effects, or another mental, neurological, or systemic disorder explains the presentation better.

MNEMONIC

Syndrome, Slope, Source, Substitutes

Syndrome identifies mild or major neurocognitive disorder; **Slope** establishes insidious progressive change; **Source** estimates the likelihood of Alzheimer's disease; **Substitutes** are competing or mixed explanations that must be actively examined.

This architecture prevents opposite errors. Premature aetiological closure labels an older patient with forgetfulness as having Alzheimer's disease before delirium, depression, medication effects, sensory limitations, or systemic illness have been considered. Endless exclusion without positive formulation lists normal investigations but never describes the characteristic course, cognitive profile, or functional consequence. The workup for reversible causes and rapidly progressive presentations is covered elsewhere in this chapter; delirium in full belongs in the Stroke, TBI, delirium and focal brain syndromes notes.

- Obtain a longitudinal history from the patient and a knowledgeable informant, anchored to change from previous ability.
- Describe affected cognitive domains and the consequences for instrumental and basic activities rather than relying on a global label.

- Look for atypical, focal, fluctuating, abrupt, stepwise, toxic, systemic, or neurological features that weaken a straightforward Alzheimer's formulation.
- State whether the evidence supports a pure aetiology, a mixed aetiology, or continuing diagnostic uncertainty.

PITFALL

Do not turn a biomarker or scan into the whole diagnosis. A positive biological signal addresses aetiology; the clinician must still define the current syndrome, functional impact, alternative contributors, capacity, risk, and care needs.

6.6 International diagnostic formulation

Under **ICD-11 6D80**, all diagnostic requirements for dementia must be met and the dementia must be presumed attributable to underlying Alzheimer disease. Attribution may draw on quantified clinical assessment or standardised cognitive or neuropsychological testing, neuroimaging, genetic testing, medical tests, family history, and clinical history. The neurological Alzheimer disease diagnosis is also assigned. This makes the distinction between the dementia syndrome and its underlying disease explicit.

The typical additional clinical description is gradual onset with progressive memory difficulty, word-finding difficulty, and mild functional impairment. The syndrome broadens as other brain regions are affected. Atypical forms express the function of the initially affected region, such as early visual-processing impairment in posterior cortical atrophy.

SUBTYPE	CLINICAL MEANING	USE
Early onset	Symptoms are present before **65 years**.	Early-onset dementia due to Alzheimer disease.
Late onset	Symptoms are present at or after the same age boundary.	Late-onset dementia due to Alzheimer disease.
Mixed with cerebrovascular disease	Alzheimer disease dementia co-exists with cerebrovascular disease.	Mixed vascular subtype.
Mixed with another nonvascular aetiology	Alzheimer disease dementia co-exists with another nonvascular aetiology.	Mixed nonvascular subtype.
Onset unknown or unspecified	Onset cannot be established reliably.	Use when onset cannot be classified more precisely.

In an examination answer, begin with the dementia syndrome, then give the presumed aetiology and the most defensible onset or mixed subtype. Do not use “mixed” as a synonym for uncertainty. It is a positive formulation that multiple pathological processes contribute materially to the cognitive syndrome.

6.7 Cross-mapping diagnostic certainty

The **DSM-5-TR** structure begins by requiring criteria for major or mild neurocognitive disorder. It then requires insidious onset and gradual progression across the affected cognitive domain or domains; major neurocognitive disorder requires impairment in at least **two domains**. The clinician next decides between probable and possible Alzheimer's disease and confirms that cerebrovascular disease, another neurodegenerative disease, substances, or another mental, neurological, or systemic disorder does not explain the disturbance better.

For major neurocognitive disorder, **probable Alzheimer's disease** is supported either by evidence of a causative mutation from family history or genetic testing, or by the complete clinical pattern: clear decline in memory and learning plus another cognitive domain, steady gradual progression without extended plateaus, and no evidence of mixed aetiology. If that standard is not met, the diagnosis is possible Alzheimer's disease.

For mild neurocognitive disorder, the certainty threshold is deliberately stricter. Probable Alzheimer's disease requires evidence of a causative mutation. Without such evidence, clear memory and learning decline, steady gradual progression without extended plateaus, and absence of mixed aetiology support **possible Alzheimer's disease**, not probable disease. The distinction reflects uncertainty about future progression at the mild stage and is a favourite source of reversed answers.

The **ICD-10-CM cross-map** in the diagnostic manual should follow the manual actually being used. Where examinations still use this older coding frame, do not substitute its coding logic for the primary international formulation. For major neurocognitive disorder due to probable or possible Alzheimer's disease, code Alzheimer disease first and then distinguish the presence or absence of behavioural disturbance; severity is recorded but not separately coded. Mild neurocognitive disorder due to Alzheimer's disease uses the mild-neurocognitive-disorder code without an additional Alzheimer disease code. The behavioural-disturbance suffixes are not reproduced here because the available source was corrupted at those characters.

6.8 Biomarker frameworks: what they add and what they do not

The approach to mild cognitive impairment due to Alzheimer's disease separates **Core Clinical Criteria**, intended for broad clinical use without highly specialised tests, from Clinical Research Criteria that incorporate biomarkers and are intended for research settings, academic centres, and clinical trials. Biomarkers are not part of the DSM mild-neurocognitive-disorder criteria described above. This is an important boundary between a clinical diagnostic manual and a research framework.

The **AT(N) framework** classifies biomarkers by the pathological process they represent. “A” denotes aggregated amyloid-beta measured through cerebrospinal-fluid amyloid measures or amyloid imaging. “T” denotes aggregated tau measured through phosphorylated tau in cerebrospinal fluid or tau imaging. “(N)” denotes neurodegeneration or neuronal injury assessed through anatomical imaging, glucose-metabolism imaging, or total tau in cerebrospinal fluid. Parentheses around N signal that neurodegeneration is not specific to Alzheimer's pathology.

Profiles are then interpreted biologically: absence of abnormal A, T, and N markers is a normal Alzheimer-biomarker profile; abnormal A alone indicates Alzheimer pathological change; abnormal A and T constitute Alzheimer's disease biologically, with N indicating whether measurable neurodegeneration is also present. An abnormal N marker without abnormal A suggests a non-Alzheimer pathological change, while abnormal A with abnormal N but normal T suggests Alzheimer pathological change with a suspected concomitant non-Alzheimer process.

This framework is intellectually useful because it separates amyloid pathology, tau pathology, and downstream injury instead of treating “a positive biomarker” as a unitary entity. Clinically, however, one must still answer what syndrome the patient has, whether functioning is impaired, whether comorbidity contributes, and whether a test result will change counselling or treatment. Cerebrospinal-fluid sampling procedure belongs in the Neurology foundations for psychiatrists notes, while dementia-specific biomarkers are examined separately in this chapter.

6.9 Management as a stage-matched care system

Management is wider than prescribing. Use a LOCUS, FOCUS, and MODUS formulation. LOCUS asks whether care can remain outpatient, requires urgent assessment, or needs inpatient management because of acute behavioural risk, delirium, medical instability, unsafe wandering, inability to meet basic needs, or collapse of the care arrangement. FOCUS identifies the immediate target and then the medium- and long-term aims. MODUS ensures that medication, psychological adaptation, environmental change, rehabilitation, social support, carer work, legal and financial planning, and palliative goals are considered together.

At every stage, disclose the formulation sensitively, assess decision-making capacity for the decision at hand, identify a reliable caregiver or support network, review driving and other safety-critical activities, simplify routines, and treat remediable contributors to distress or poor performance. Preserve autonomy where it remains safe. The care plan should anticipate the next likely transition rather than waiting for a crisis to reveal that supervision, respite, or a different living arrangement is already needed.

For cognitive symptoms, cholinesterase inhibitors and memantine are the established symptomatic drug classes, selected according to severity, tolerability, comorbidity, and the applicable regulatory framework. Their dosing, titration, adverse effects, interactions, licensing bands, and stopping decisions are taught in the dedicated cognitive-enhancer section rather than duplicated here.

Disease-modifying anti-amyloid therapy requires separate assessment, biological confirmation, and monitoring. Lecanemab, donanemab, and aducanumab differ in evidence, regulatory history, administration, monitoring burden, and risk. Those protocols, efficacy estimates, and amyloid-related imaging abnormalities belong in the dedicated anti-amyloid section. They should never displace symptomatic treatment, vascular and general health care, psychosocial intervention, or support for the caregiver.

Behavioural and psychological symptoms require a formulation before a prescription: identify pain, fear, overstimulation, unmet needs, communication failure, medication effects,

environmental triggers, and superimposed medical illness. Begin with targeted non-pharmacological and environmental measures whenever the situation allows. Antipsychotic use is exceptional, risk-based, time-limited, and reviewed. Dementia-specific prescribing, together with stroke and mortality risk and the wider management of behavioural symptoms, is covered in the dedicated behavioural-symptom section. General antipsychotic pharmacology belongs in the Schizophrenia: management, antipsychotics and adverse effects notes.

In advanced disease, FOCUS shifts toward comfort, dignity, prevention of avoidable injury, feeding and swallowing decisions, pressure care, communication through familiar routines, and support for family grief and burden. Palliative principles can coexist with active treatment of distress. The final care plan should be explicit about who will respond to deterioration, where care is preferred, which interventions remain proportionate, and how carers can obtain urgent help.

Diagnose the syndrome, demonstrate the progressive course, grade the certainty of Alzheimer's aetiology, exclude a better or mixed explanation, and match the care plan to function and risk.

7 · Dementia with Lewy bodies: core and supportive features and the one-year rule

The marks lie in classification, not in recognition alone. A patient with cognitive decline, vivid visual experiences, changing attention, dream enactment and parkinsonism may suggest dementia with Lewy bodies immediately. The higher-quality answer goes further. It establishes the essential dementia syndrome, separates core clinical features from supportive findings, distinguishes indicative from supportive biomarkers, applies the probable or possible threshold correctly, and then uses the temporal sequence of cognitive and motor symptoms to distinguish dementia with Lewy bodies from dementia arising in established Parkinson disease.

The central discipline is to keep different diagnostic frameworks apart. The current consortium model gives REM sleep behaviour disorder the same core status as the other core clinical features. The **DSM-5-TR framework** retains an older split in which REM sleep behaviour disorder is suggestive rather than core. Each framework can appear in examinations and clinical records. A safe answer therefore names the framework before counting features.

50-75%	ALMOST 85%	60-92%	AS HIGH AS 77%
FLUCTUATING COGNITION	COMPLEX VISUAL HALLUCINATIONS	SPONTANEOUS PARKINSONISM MAY OCCUR	REM SLEEP BEHAVIOUR DISORDER

7.1 Start with the essential syndrome

The essential feature is dementia: progressive cognitive decline severe enough to interfere with usual social or occupational functioning. This requirement prevents a collection of characteristic symptoms from being mistaken for a dementia diagnosis when everyday independence is not meaningfully affected. The core features refine the aetiologic diagnosis only after the cognitive syndrome has been established.

The cognitive profile may not begin as a classic amnestic presentation. Prominent or persistent memory impairment may be absent early and usually becomes evident as the disorder progresses. Attention, executive function and visuospatial ability may be especially impaired. Clinically, this means that a patient who struggles with spatial judgement, planning, mental tracking or sustained attention may have disabling cognitive decline even when relatives do not initially describe repetitive forgetting as the dominant problem.

This profile also explains why a brief memory-focused history can miss the syndrome. The interview should reconstruct performance across tasks: navigating familiar spaces, organising multistep activities, following a conversation through variable alertness, judging visual relationships and adapting when a rule changes. Functional decline remains the anchor. A low test performance without corroborated decline is not interchangeable with dementia, while apparently adequate recall does not exclude a disorder whose early burden lies elsewhere.

RULE **Establish dementia before classifying the Lewy-body pattern.** Core features and biomarkers determine diagnostic probability only after progressive, functionally significant cognitive decline has been demonstrated.

7.2 The four core clinical features

The current McKeith 2017 **DLB Consortium model** recognises **four** core clinical features: fluctuating cognition, recurrent well-formed visual hallucinations, spontaneous parkinsonism and REM sleep behaviour disorder. The accompanying figure provides the visual spine. The candidate should be able not only to name these features, but also to identify what makes each one diagnostically characteristic.

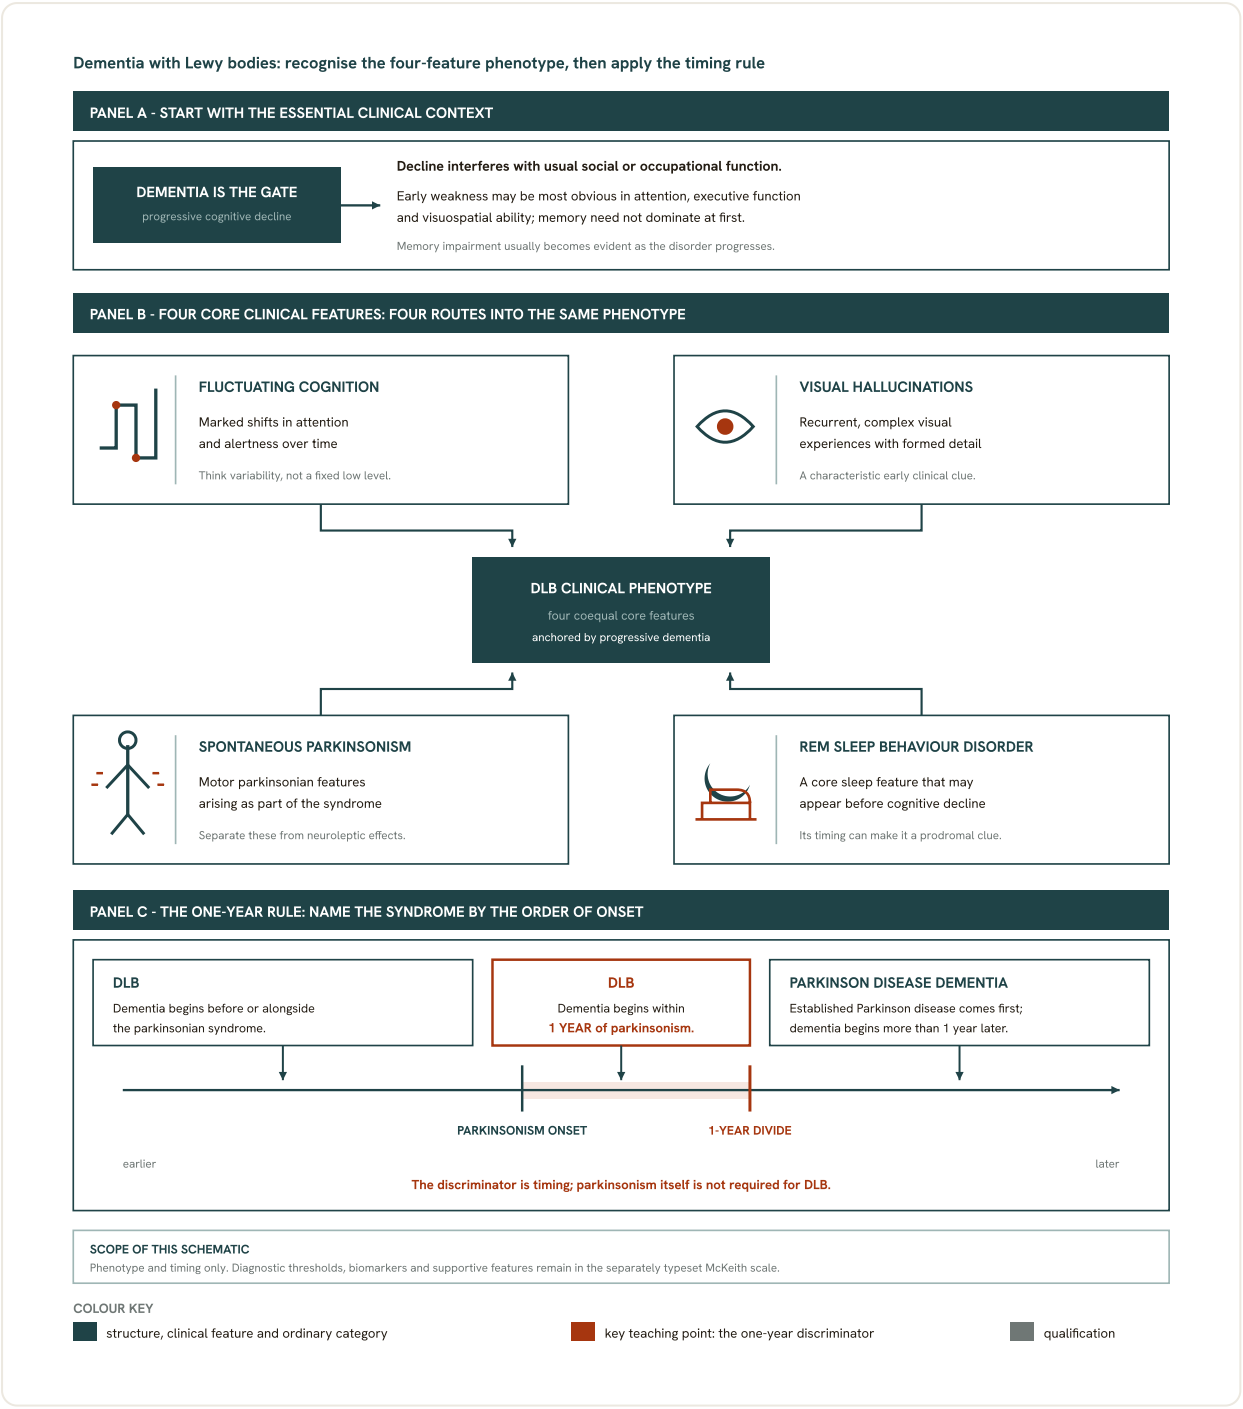


Fig 32.4 **Dementia with Lewy bodies: four core clinical features and the one-year rule.** Progressive dementia provides the clinical context; fluctuating cognition, formed recurrent visual hallucinations, spontaneous parkinsonism and REM sleep behaviour disorder form the core phenotype. The relative timing of dementia and parkinsonism separates dementia with Lewy bodies from Parkinson disease dementia. Adapted conceptually from the DLB Consortium fourth consensus report (McKeith et al., *Neurology* 2017;89:88-100).

Fluctuating cognition refers to pronounced variation in attention and alertness. It is not simply a good morning followed by a tired evening, nor the understandable variability produced by effort, anxiety or poor sleep. The diagnostically useful history describes conspicuous changes in engagement, coherence or responsiveness over time. Collateral history is crucial because the patient may not recognise the contrast, and the clinic encounter may capture either the clearer or the impaired state.

Recurrent complex visual hallucinations are typically vivid, colourful, well formed, three-dimensional and silent, often depicting people, children or animals. They are usually recognisable scenes or figures rather than vague visual distortion. Their recurrence, detail and early presence support the pattern. This phenomenology should be elicited neutrally: ask what the person sees, how formed the image is, whether it moves or speaks, how often it returns and how convinced the person is that it is present. In dementia with Lewy bodies these hallucinations tend to occur early, whereas hallucinations in Alzheimer disease are generally a later complication.

Spontaneous parkinsonism means motor parkinsonian features arising as part of the illness rather than as an adverse effect of dopamine-blocking treatment. The distinction is essential because neuroleptic-induced extrapyramidal signs can falsely create an apparently core feature. Parkinsonism is not mandatory: up to 25% of people with probable neurocognitive disorder with Lewy bodies may never develop extrapyramidal signs. Absence of the motor syndrome therefore does not cancel a diagnosis supported by the remaining clinical and biomarker pattern.

REM sleep behaviour disorder is dream enactment associated with loss of normal REM sleep atonia. It may precede cognitive decline and often begins several years before the dementia syndrome; it is likely to be a common prodromal feature. Ask the bed partner about purposeful or defensive movements during dreams, not merely restless sleep. When the history is uncertain, polysomnographic demonstration of REM sleep without atonia serves as an indicative biomarker rather than relying on description alone.

MNEMONIC

F-V-P-R

Fluctuating cognition, **V**isual hallucinations, spontaneous **P**arkinsonism and **R**EM sleep behaviour disorder form the current consortium core. The mnemonic is a recall frame; the marks come from describing the characteristic form of each feature.

7.3 Supportive clinical features: persuasive, but unweighted

Supportive clinical features strengthen pattern recognition and often dominate practical risk, but they carry no formal diagnostic weighting in the consortium algorithm. This distinction is easy to lose because some supportive findings are striking. They matter to assessment and care, yet they cannot be added together as if they were core features or indicative biomarkers.

- Severe sensitivity to antipsychotic agents, postural instability, repeated falls, syncope and other transient episodes of unresponsiveness are clinically consequential supportive features.
- Severe autonomic dysfunction broadens the history beyond cognition and movement; the symptom should be linked to functional risk rather than recorded as an isolated checklist item.
- Hypersomnia and hyposmia may support the overall syndrome but do not independently set its diagnostic probability.
- Hallucinations in other modalities, systematised delusions, apathy, anxiety and depression describe the neuropsychiatric breadth of the illness without acquiring core-feature status.

The absence of weighting does not mean absence of importance. Repeated falls, unresponsiveness and autonomic dysfunction can alter the urgency and setting of assessment. Antipsychotic sensitivity can change prescribing safety. The detailed management of neuroleptic sensitivity belongs to the dedicated section on dementia with Lewy bodies management, while general antipsychotic pharmacology belongs to the Schizophrenia: management, antipsychotics and adverse effects notes. Here, the diagnostic lesson is narrower: supportive findings enrich the phenotype but do not cross the threshold by arithmetic.

PITFALL

Do not count several supportive symptoms as equivalent to a core feature. A patient may have falls, autonomic symptoms, depression and delusions, yet the probable or possible category still depends on the specified combination of core features and indicative biomarkers.

7.4 Indicative and supportive biomarkers occupy different levels

Indicative biomarkers are reduced dopamine-transporter uptake in the basal ganglia on SPECT or PET, low cardiac uptake on iodine-labelled MIBG myocardial scintigraphy, and polysomnographic confirmation of REM sleep without atonia. They enter the probable or possible algorithm directly. They should be ordered to resolve a diagnostic question, not as decorative additions after the clinical pattern is already secure.

Dopamine-transporter imaging is especially useful when the motor examination does not clearly establish parkinsonism. In a meta-analysis of four studies involving 419 people, the cited SPECT method had **86.5% sensitivity and 93.6% specificity** for diagnosing dementia with Lewy bodies. The same source notes that the test is unnecessary once parkinsonism is well established clinically. This is a practical example of diagnostic restraint: a biomarker is most useful when it can materially change confidence, not merely repeat an obvious bedside finding.

For myocardial MIBG scintigraphy, a meta-analysis of eight studies involving 346 people reported 98% sensitivity and 94% specificity for differentiating dementia with Lewy bodies from other dementias. Such figures describe performance in the cited evidence base, not a substitute for examining the patient, checking competing explanations or applying the clinical criteria. Imaging physics and modality principles are covered in the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

Supportive biomarkers include relative preservation of medial temporal structures on CT or MRI; generalised low perfusion or metabolism with reduced occipital activity, with or without a cingulate island sign on FDG-PET; and prominent posterior slow-wave EEG activity with periodic fluctuation in the pre-alpha or theta range. These findings support the pattern but carry no diagnostic weighting. The conceptual boundary mirrors the clinical one: “supportive” means coherent with the diagnosis, not numerically exchangeable with an indicative biomarker.

DIAGNOSTIC LAYER	WHAT BELONGS HERE	HOW IT IS USED
Essential syndrome	Progressive cognitive decline that interferes with social or occupational functioning	Required before probable or possible DLB is considered
Core clinical features	Fluctuation, recurrent well-formed visual hallucinations, spontaneous parkinsonism and REM sleep behaviour disorder	Directly determine diagnostic probability
Indicative biomarkers	Reduced basal-ganglia dopamine-transporter uptake, low myocardial MIBG uptake, or REM sleep without atonia	Directly enter the criteria in combination with the clinical pattern
Supportive clinical features	Sensitivity, falls, unresponsiveness, autonomic and sleep symptoms, other hallucinations, delusions and affective symptoms	No diagnostic weighting
Supportive biomarkers	Medial temporal preservation, occipital functional-imaging change, cingulate island sign or posterior slow-wave EEG activity	Support the pattern without diagnostic weighting
Factors against the diagnosis	Another physical or brain disorder sufficient to explain the picture, or parkinsonism first appearing only with severe dementia	Reduce diagnostic likelihood and force reconsideration

7.5 Converting evidence into probable or possible DLB

The consortium threshold is deliberately asymmetric. **Two** core clinical features are sufficient for probable dementia with Lewy bodies, whether or not an indicative biomarker is present. Probable DLB can also be diagnosed when one core clinical feature is accompanied by one or more indicative biomarkers. Indicative biomarkers alone cannot establish probable DLB. This last clause matters: strong biological support without a core clinical expression remains below the probable threshold.

One core clinical feature is sufficient for possible dementia with Lewy bodies when indicative biomarker evidence is absent. Possible DLB can also be diagnosed when one or more indicative biomarkers are present but no core clinical feature is present. The category is inherently less stable: some patients later fulfil probable DLB, while others evolve toward Alzheimer disease or another disorder. “Possible” should therefore prompt longitudinal review rather than false certainty.

The algorithm is easiest to apply in sequence. Count only core clinical features, then ask whether an indicative biomarker is present. Do not count supportive clinical findings, supportive biomarkers or the mere clinical impression of a “Lewy-like” illness. This ordering prevents

criterion inflation and keeps the probable category tied to either a sufficiently rich clinical phenotype or a clinical-biomarker convergence.

GOOD TO REMEMBER

When writing the diagnosis, show the evidence trail concisely: establish dementia, name the core feature or features, name any indicative biomarker, and then state probable or possible DLB. This makes a counting error visible before it reaches the conclusion.

7.6 The criteria-set trap: consortium versus the diagnostic manual

The consortium and DSM-5-TR frameworks use different taxonomies. In the consortium framework, REM sleep behaviour disorder is one of four core clinical features. In the manual framework, there are three core features: cognitive fluctuation, recurrent well-formed visual hallucinations and spontaneous parkinsonism arising after cognitive decline. REM sleep behaviour disorder and severe neuroleptic sensitivity are suggestive features there.

The resulting thresholds also differ in language. Under DSM-5-TR, probable disorder requires two core features, or one or more core features together with a suggestive feature; possible disorder requires only one core feature, or suggestive features without any core feature. Under the current consortium model, the probability calculation instead combines its four core clinical features with indicative biomarkers. The same patient can therefore be described through legitimate but non-identical structures.

TRAP **Do not answer "How many core features?" until the criteria set is named.** The current consortium report has four core clinical features; DSM-5-TR has three core features and places REM sleep behaviour disorder among suggestive features. Silently merging them creates a hybrid system that belongs to neither source.

This is also why "neuroleptic sensitivity is a core feature" is an unsafe unqualified statement. It is supportive in the consortium clinical list and suggestive in the manual framework, but not core in either of the supplied schemes. Conversely, REM sleep behaviour disorder changed status in the updated consortium approach while the manual framework retained its earlier position. Attribution, not memorised confidence, resolves the apparent contradiction.

7.7 The one-year rule: sequence, label and limitation

The distinction rests on when dementia emerges relative to the motor disorder. Dementia with Lewy bodies is diagnosed when dementia occurs before or concurrently with parkinsonism. Parkinson disease dementia is the term used when dementia occurs in the context of well-established Parkinson disease. For research comparisons, the existing **one-year rule** remains recommended. In ordinary clinical settings, the label most appropriate to the presentation may be used, and a broader term such as Lewy body disease can sometimes communicate the continuum more honestly.

The DSM-5-TR formulation states the other side of the boundary: for major neurocognitive disorder to be attributed to Parkinson disease, the Parkinson disease diagnosis must have been

present for **at least one year** before cognitive decline reaches the level of major neurocognitive disorder. Cognitive symptoms in neurocognitive disorder with Lewy bodies may begin before, with or even without parkinsonism. Thus, the rule is not "motor symptoms at any time means Parkinson disease dementia." It is a temporal convention anchored to established motor disease preceding the dementia syndrome.

CLINICAL SEQUENCE	PREFERRED LABEL	INTERPRETIVE POINT
Dementia before parkinsonism	Dementia with Lewy bodies	Cognitive syndrome leads the motor syndrome
Dementia and parkinsonism begin concurrently	Dementia with Lewy bodies	Concurrent onset remains on the DLB side of the convention
Established Parkinson disease precedes major cognitive decline by at least one year	Neurocognitive disorder due to Parkinson disease or Parkinson disease dementia	Motor disease clearly precedes the dementia syndrome
Clinical presentation does not fit a clean temporal split	Use the clinically appropriate term; Lewy body disease may be helpful	The strict interval is primarily retained for research distinction

The convention has acknowledged limitations. It may arbitrarily divide a continuum of presentations arising from underlying Lewy body disease, and it offers little treatment guidance by itself. An additional Parkinson disease expert-consensus view permits clinician choice between the two diagnoses when major neurocognitive disorder begins before or within twelve months of the motor diagnosis. The existence of this alternative is precisely why the candidate should state which convention is being applied rather than present the boundary as a biological discontinuity.

Parkinson disease neuropsychiatry without dementia, including psychosis outside the dementia syndrome, belongs to the Neuropsychiatry of systemic and neurological disease notes. The present distinction is limited to the dementia lens: clinical sequence determines nomenclature, while the broader motor and neuropsychiatric disorder requires its own assessment.

7.8 Differentiation from Alzheimer disease and diagnostic restraint

Alzheimer disease is the most common clinical misdiagnosis of dementia with Lewy bodies. The neuropsychological contrast is relative rather than absolute. Compared with Alzheimer disease, people with DLB may perform worse on attention, working memory, figure copying, spatial judgement and set shifting. Visuospatial and executive impairment are often greater, while episodic and verbal memory impairment may be less severe. Better recognition than free recall suggests a retrieval problem rather than complete loss of stored information.

The clinical features refine this profile. Early recurrent, detailed visual hallucinations, pronounced cognitive fluctuation, spontaneous parkinsonism and REM sleep behaviour disorder shift the pattern toward DLB. The presence of at least two early core features among fluctuation, prominent visual hallucinations and parkinsonism has high specificity for DLB in the cited account. Functional imaging may add reduced dopamine-transporter uptake in the basal ganglia

or occipital hypoperfusion or hypometabolism. A positive dopamine-transporter scan is not entirely exclusive to DLB, while associated occipital hypometabolism favours DLB over frontotemporal dementia.

Even a convincing clinical distinction does not imply mutually exclusive pathology. A person with a typical insidious amnestic Alzheimer presentation may have Lewy-body pathology at autopsy, and a patient may meet criteria for both disorders. The correct conclusion is therefore probabilistic: identify the dominant clinical syndrome, acknowledge convergent biomarker evidence and remain open to co-pathology. An isolated hallucination, scan result or extrapyramidal sign should not be allowed to erase the rest of the phenotype.

TRAP **Do not quote a prevalence estimate as though all sources measured the same population.** Clinically diagnosed DLB averages 4.2% of dementias in community studies and 7.5% in clinical samples in the cited population-framed estimates, whereas Lewy-body pathology at autopsy is reported in 20-35% of dementia cases in US and UK studies. Other textbooks give higher general clinical estimates. These are different denominators and ascertainment methods, not numbers to average.

The final examination answer should therefore preserve hierarchy and attribution. Begin with functionally significant progressive cognitive decline. Describe the current consortium core, then the supportive clinical features. Separate indicative biomarkers, which can change probable or possible status, from supportive biomarkers, which cannot. State the probability threshold. Flag the alternative manual taxonomy. Finally, apply the temporal rule to distinguish DLB from Parkinson disease dementia and note that the interval is a practical convention imposed on a clinical continuum.

Memorise the sequence: establish dementia; identify the four consortium core features; remember that supportive findings carry no weight; apply the rule that probable DLB requires two core features or one core plus an indicative biomarker; and recognise that established Parkinson disease preceding major cognitive decline by at least one year favours Parkinson disease dementia.

8 · Dementia with Lewy bodies: neuroleptic sensitivity, safe prescribing and management

The treatment paradox. Dementia with Lewy bodies combines cognitive, psychiatric, motor, sleep and autonomic problems, yet the drug most likely to be reached for during frightening hallucinations or agitation may cause the most dangerous deterioration. **Severe neuroleptic sensitivity** can present with an acute onset or marked worsening of parkinsonism and impaired consciousness after even a low dose of a first-generation or second-generation antipsychotic. Tolerance of a previous exposure does not exclude the diagnosis.

This is therefore not merely a psychosis-treatment topic. It is a diagnostic-safety and prescribing topic. **Up to 50%** of people with neurocognitive disorder with Lewy bodies may show severe neuroleptic sensitivity, so recognition of the syndrome changes the entire management plan. The

high-yield answer begins by deciding whether the experience requires drug treatment at all, removes reversible aggravators, uses cognitive-enhancing treatment as part of behavioural management, and reserves dopamine-blocking treatment for exceptional situations.

RULE **In suspected or established DLB, an antipsychotic is an exception, not the reflex response to a hallucination.** Treat the patient's distress and risk, not the mere presence of a perceptual symptom.

8.1 Management frame: LOCUS, FOCUS and MODUS

LOCUS determines how intensively the patient can be assessed and observed. Stable cognitive or behavioural symptoms usually belong in outpatient care with a reliable caregiver, medication reconciliation and planned review. Inpatient care is appropriate when behaviour cannot be safely contained, consciousness or mobility has deteriorated, oral intake is compromised, or close observation is required while potentially offending medicines are withdrawn. Emergency care is required for an abrupt severe reaction after antipsychotic exposure or for immediate danger that cannot be managed in the current setting.

FOCUS should be stated as an observable target rather than as "treat DLB." Acute targets are safety, severe distress, impaired consciousness and sudden motor worsening. Short-term targets include hallucinations that frighten the patient, delusional behaviour that creates danger, sleep-related injury, depressive symptoms, falls and caregiver exhaustion. Mid-term targets are cognition, daily function, mobility and a sustainable medication burden. Long-term targets are quality of life, participation, caregiver capability and avoidance of preventable iatrogenic decline.

MODUS includes environmental and social measures, caregiver education, cognitive and behavioural treatment, carefully selected medication, rehabilitation and, in narrowly selected severe presentations, a procedural treatment. The patient and caregiver should understand that symptoms may fluctuate, that a calm hallucination does not automatically need suppression, and that every new medicine must earn its place against sedation, postural symptoms, confusion, reduced mobility and polypharmacy.

The governing prescription style is **geriatric psychopharmacology**: initiate at a very low dose, titrate slowly, monitor periodically, reduce the dose at the earliest reasonable opportunity and minimise polypharmacy. In DLB this is not ceremonial caution. It is the practical response to unusually high sensitivity to drug-related adverse effects.

AVOID	START	TITRATE	REVIEW
ROUTINE BLOCKADE	VERY LOW	SLOWLY	EARLY

8.2 Recognising and responding to neuroleptic sensitivity

The core clinical pattern is temporal: a dopamine-blocking drug is introduced or increased, followed by a disproportionate acute deterioration in parkinsonism, mobility, alertness or consciousness. The response may occur despite a dose that appears modest by ordinary psychiatric practice. This makes the medication timeline as important as the mental-state

description. Ask what was given, why it was given, what changed afterwards, and whether the change was motor, cognitive, behavioural or global.

First-generation antipsychotics should generally be avoided. Substituting a second-generation agent does not erase the risk because tolerance may also be low and the evidence guiding such use is limited. If a severe reaction is suspected, withhold further exposure, move the patient to an appropriate medical LOCUS, assess consciousness and motor function, look for concurrent illness or medication effects, and provide supportive monitoring. Rechallenge should never be casual merely because psychotic symptoms persist.

CLINICAL QUESTION	INTERPRETATION	IMMEDIATE MANAGEMENT IMPLICATION
Is the hallucination distressing or dangerous?	A well-tolerated symptom may not justify drug risk	Observe, explain and adjust the environment
Was an antipsychotic recently introduced or increased?	Creates a direct temporal suspicion	Withhold further exposure and reassess urgently
Have parkinsonism and consciousness worsened together?	Strongly supports a severe sensitivity response	Escalate the LOCUS and monitor closely
Is another precipitant plausible?	Illness, pain, dehydration and medication burden may amplify symptoms	Search for and correct the precipitant
Did the patient previously tolerate an antipsychotic?	Prior tolerance does not rule out DLB	Keep the diagnosis and sensitivity warning active

PITFALL

Do not label abrupt post-prescription immobility and reduced responsiveness as “progressing dementia” without reconstructing the medication timeline. The error converts an adverse reaction into an excuse for further prescribing.

Altered awareness may also indicate delirium superimposed on dementia. Its full assessment and treatment belong in the Stroke, TBI, delirium and focal brain syndromes notes. General antipsychotic pharmacology and the broader extrapyramidal and metabolic adverse-effect teaching belong in the Schizophrenia: management, antipsychotics and adverse effects notes.

8.3 The sequence before any antipsychotic

Start with a formulation, not a prescription. Clarify whether the target is hallucination, delusion, agitation, fear, aggression, sleep disruption or caregiver distress. Establish whether the patient is suffering or at risk. Review pain, infection, constipation, urinary symptoms, hydration, sensory impairment, sleep disruption, unfamiliar surroundings and recent medication changes. These checks do not deny the neuropsychiatric symptom; they prevent a modifiable driver from being mistaken for an indication for dopamine blockade.

The recommended **therapeutic sequence** is clinically coherent. If distressing hallucinations or other neuropsychiatric symptoms occur in a person receiving anti-parkinsonian medication, first

consider whether that regimen can be reduced without unacceptable motor loss. Then consider a trial of an acetylcholinesterase inhibitor. Only if symptoms remain sufficiently severe should a low-dose dopamine antagonist enter the discussion, because these drugs are poorly tolerated and their supporting evidence is inadequate.

MNEMONIC

PAUSE before blockade

Pin down the target and precipitant; **A**ssess distress, danger and anti-parkinsonian medication; **U**se environmental and caregiver measures; **S**tart cognitive-enhancing treatment when indicated; **E**xceptionally consider an antipsychotic with explicit monitoring.

Non-drug care should be concrete. Improve lighting and visual cues, reduce confusing shadows and overstimulation, preserve familiar routines, check hearing and vision, and coach caregivers to acknowledge distress without arguing about the content of the experience. Increasing arousal and attention through social interaction or environmental novelty may lessen visual hallucinations and fluctuations. A written plan should state which behaviour warrants contact, which change warrants emergency review, and which medicines should not be given reflexively.

The purpose of sequencing is not therapeutic delay. It matches treatment intensity to symptom burden while preserving cognition, mobility and consciousness. When danger is immediate, containment and urgent assessment proceed in parallel; when danger is absent, time spent on formulation is often the safest intervention.

8.4 If an antipsychotic is unavoidable

An unavoidable prescription requires a documented indication: severe distress, dangerous behaviour or psychosis that has not responded to safer measures. Record the baseline motor state, alertness, postural symptoms, oral intake and functional ability; discuss the sensitivity risk with the caregiver; choose the least burdensome option supported by the available evidence; use the lowest practical exposure; and arrange early review. Continued treatment should depend on demonstrable benefit, not on the inertia of a discharge list.

Write the monitoring contract before exposure. It should contain:

- a named target behaviour, with concrete examples of the distress or danger being treated;
- a baseline description of movement, consciousness, posture, intake and the patient's usual fluctuation;
- caregiver instructions to report sudden immobility, reduced responsiveness, falls, excessive sedation or functional loss; and
- a prespecified review decision: continue only for clear benefit, otherwise reduce or withdraw promptly.

This contract prevents common distortions. A transient quietening caused by sedation must not be mistaken for treatment of psychosis, and a naturally fluctuating good period must not be

credited automatically to the drug. The relevant outcome is safer, less distressing behaviour without a disproportionate cost to consciousness, mobility, cognition or daily function.

The evidence is limited and sometimes internally uncomfortable:

AGENT	WHAT THE DLB EVIDENCE SAYS	PRESCRIBING INTERPRETATION
Quetiapine	An open study suggested benefit in some patients but somnolence and orthostatic hypotension caused discontinuation; a controlled study found tolerability without superiority to placebo	Tolerability in the controlled study does not establish efficacy
Risperidone	Poor tolerability, with worsening of cognition and neuropsychiatric symptoms reported	Does not offer a reassuring DLB profile
Olanzapine	Generally poorly tolerated, although a small trial found **5 mg/day** reduced hallucination and delusion severity	This is a dose-specific trial finding, not a general recommended DLB dose
Clozapine	No DLB-specific efficacy study; **6.25-50 mg/day** is extrapolated from psychosis in Parkinson disease dementia	Any consideration must acknowledge that the dose and efficacy inference are not derived from a DLB trial

The Parkinson-disease-psychosis-without-dementia algorithm, including detailed use of pimavanserin, belongs in the Neuropsychiatry of systemic and neurological disease notes. The distinction matters: importing an algorithm from a neighbouring disorder without declaring the extrapolation turns uncertainty into false authority.

TRAP **"Use an atypical antipsychotic because it is safer" is not an adequate DLB answer.** Second-generation status does not neutralise neuroleptic sensitivity; the answer must state avoidance, exceptional use, very low initiation, slow titration, explicit monitoring and early withdrawal when benefit is absent.

GOOD TO REMEMBER

Put "possible severe neuroleptic sensitivity" prominently in the case record and caregiver plan. A correct diagnosis that is invisible during an emergency admission does not protect the patient.

8.5 Cognition and behaviour: cholinesterase inhibitors and memantine

Cholinesterase inhibitors are the mainstay of DLB treatment, with the strongest evidence for donepezil and rivastigmine. Their value is broader than a cognitive-score change: beneficial effects have been reported for cognition, hallucinations, delusions, activities of daily living and global clinical status. Adverse effects include nausea, vomiting, anorexia and somnolence, with adverse events particularly prominent with rivastigmine; its patch may reduce some gastrointestinal difficulty associated with oral treatment. Evidence for galantamine is more limited.

The **NICE sequencing for DLB cognition** is a useful examination and clinical spine. Offer donepezil or rivastigmine in mild to moderate illness; consider galantamine only when both are not tolerated. Donepezil or rivastigmine may also be considered in severe illness. Consider memantine when acetylcholinesterase inhibitors are not tolerated or are contraindicated. The sequence identifies choice and place in therapy; it does not supply a DLB-specific dose in the evidence available here.

Trial evidence supports the hierarchy without making it absolute. Donepezil produced a pooled cognitive advantage of **nearly 2 MMSE points**, with maintained cognitive benefit and a greater likelihood of avoiding deterioration on global impression. Rivastigmine showed a smaller cognitive change that was not statistically significant. These data should support a monitored therapeutic trial, not substitute a score for the patient's function. Ask whether attention, engagement, hallucinations, self-care and caregiver burden have changed, and document tolerability alongside benefit.

Memantine has mixed evidence in DLB. Studies using mixed DLB and Parkinson disease dementia samples found small but statistically significant benefit in DLB subgroup analyses, while individual reports describe possible improvement in quality of life, global status, goal attainment, caregiver burden, attention and sleep behaviour. The correct conclusion is modest uncertainty: memantine has a defined place when cholinesterase inhibitors cannot be used, but it should not be presented as uniformly effective.

Before initiation, agree on target outcomes and a review method. During follow-up, separate adverse effects from disease fluctuation, and ask caregivers for examples rather than a global "better" or "worse." Formal cognitive instruments may support longitudinal observation, but their general administration and scoring belong in the Psychometry, Assessment and Descriptive Psychopathology notes.

A useful follow-up review triangulates the patient's account, caregiver observation and direct examination. Cognitive benefit may appear as steadier attention, more coherent conversation, greater participation in self-care or fewer distressing perceptual experiences rather than as a dramatic test-score change. Conversely, nausea, poor appetite, daytime sleepiness, worsening postural symptoms or reduced engagement may erase a modest benefit. Document both sides of the ledger. If the disease is fluctuating, repeat observation across ordinary settings before concluding that treatment has failed, unless an adverse effect demands earlier action.

8.6 Motor, sleep, autonomic and affective targets

DLB management is symptom-targeted because improving one domain can destabilise another. Anti-parkinsonian medication may aid mobility yet aggravate hallucinations; reducing it may ease psychosis yet impair movement. The practical method is negotiated prioritisation: identify which loss is currently most disabling, alter contributors sequentially where possible, and judge the whole patient rather than a single symptom scale. A **phase 2** placebo-controlled trial found that adjunctive zonisamide improved parkinsonism without worsening cognition or psychiatric

symptoms, but adverse effects increased with dosage, particularly weight loss and reduced appetite. No DLB-specific dose is supplied by that evidence.

REM sleep behaviour disorder requires an injury-prevention plan as well as medication consideration. Make the sleeping environment safer and ask about dream enactment, falls from bed, injury to the patient or bed partner, nocturnal confusion and daytime sedation. Rivastigmine or memantine may reduce episode frequency. Melatonin and low-dose clonazepam are recognised options, while clonazepam requires particular caution in older adults because it is long acting. The available DLB passages do not establish a specific dose for either agent.

For depression and anxiety, assess severity, suicide risk, medical contributors and medication burden, then use psychological, social and pharmacological options cautiously. **Citalopram** is the only antidepressant investigated in a controlled DLB trial; it showed no benefit and was poorly tolerated. Other controlled DLB trials of selective serotonin reuptake inhibitors or serotonin-noradrenaline reuptake inhibitors are lacking, and medicines with anticholinergic adverse effects should be avoided. Postural hypotension, urinary symptoms, restless legs and excessive daytime sleepiness should be treated as explicit targets rather than folded into "dementia." Each intervention must be checked for its effect on cognition, falls, alertness and interaction with the rest of the regimen.

This multidomain review is where polypharmacy grows unless someone owns the whole list. At every visit ask: What is each medicine targeting? Is benefit visible? Is the symptom still active? Can the dose be reduced or the drug stopped? The aim is not a medicine for every symptom; it is the smallest coherent regimen compatible with safety and function.

8.7 Psychological, social, rehabilitative and procedural care

Non-pharmacological care is not an appendix to drug treatment. Person-centred care, physical exercise, cognitive training, social interaction and cognitive behavioural approaches may help, although the evidence remains limited. Environmental novelty and purposeful interaction may improve arousal and attention and thereby reduce hallucinations or functional fluctuation. Occupational and physiotherapy input can translate these principles into safer mobility, preserved routine and meaningful activity.

Caregiver work has direct clinical value. Explain fluctuations, hallucinations and medication sensitivity; distinguish reassurance from confrontation; plan responses to nocturnal behaviour; reduce avoidable environmental confusion; and provide a route for urgent help. Review caregiver sleep, strain and ability to supervise medicines. Social planning should include fall risk, home safety, driving or other hazardous activities, adherence, advance preferences and the realistic capacity of the household. These discussions should be revisited as cognition and function change.

Transitions between home, emergency care and the ward are high-risk moments because the clinician who receives the patient may see agitation without knowing the usual motor state or the sensitivity history. Send a compact care summary with the diagnosis, baseline cognition and mobility, medicines that previously caused deterioration, current targets, caregiver contact and

agreed escalation plan. On admission, reconcile the actual medicines taken rather than copying an old list. At discharge, remove discontinued drugs from the active list, explain every remaining change to the caregiver and specify who will review efficacy and adverse effects. Continuity is itself a safety intervention when the syndrome fluctuates and multiple specialties are involved.

Procedural treatment is reserved for selected severe illness. Electroconvulsive therapy may be considered when severe depression and psychosis coexist, but the DLB evidence is limited to case reports and a small case series. Repetitive transcranial magnetic stimulation may also be helpful, with similarly limited certainty. The decision therefore requires a clear target syndrome, appraisal of medical and cognitive vulnerability, caregiver involvement and follow-up that tests whether meaningful benefit occurred.

> In DLB, protect the patient from iatrogenic dopamine blockade before treating distressing symptoms through a staged, target-based plan that preserves cognition, consciousness and mobility.

9 · Vascular dementia: subtypes, clinical features, criteria and the Hachinski score

Why it matters clinically. Vascular dementia is not a unitary stereotyped syndrome. It is a clinically defined neurocognitive disorder in which cerebrovascular disease is sufficient to explain the cognitive decline. The marks lie in showing that the candidate can connect phenotype, course and vascular evidence, rather than merely noticing vascular risk factors or an abnormal scan. Across major sources, vascular disease is the **second** most common cause of dementia after Alzheimer disease. Its importance therefore comes from frequency as well as diagnostic complexity.

The modern term **vascular neurocognitive disorder** includes major and mild syndromes, while the traditional term vascular dementia usually implies cognitive impairment severe enough to interfere substantially with independence. This section uses both where the source framework requires it. The broader vascular cognitive impairment spectrum, post-stroke cognitive syndromes and detailed strategic-infarct neuropsychiatry belong in the Stroke, TBI, delirium and focal brain syndromes notes. Here the task is narrower: recognise vascular dementia, classify its principal forms, describe its characteristic clinical profiles, apply diagnostic criteria and interpret the ischaemic scale without overclaiming what it proves.

RULE **Vascular disease must be explanatory, not merely present.** A diagnosis rests on a coherent relationship between cognitive decline and cerebrovascular disease, supported by the course, examination and appropriate imaging.

9.1 The diagnostic idea: syndrome plus sufficient vascular cause

A short sequence of questions organises the assessment. Is there a major or mild neurocognitive syndrome? Is cerebrovascular disease demonstrated clinically, on examination or on neuroimaging? Is that disease sufficient, in distribution and timing, to account for the cognitive

deficits? This sequence prevents the common shortcut of moving directly from a vascular risk factor to a dementia label.

The relationship may announce itself through cognitive decline close to a cerebrovascular event. It may instead appear as a pattern dominated by slowed information processing, impaired complex attention and frontal-executive dysfunction. Neither route eliminates the need to consider alternative brain disease or systemic disorder. Conversely, absence of a dramatic stepwise history does not rule out a vascular disorder, because gradual and slowly progressive presentations occur with small-vessel disease affecting white matter, basal ganglia or thalamic structures.

Temporal	Executive	Focal	Imaging
EVENT LINK	COGNITIVE PROFILE	CLINICAL EVIDENCE	PARENCHYMAL INJURY

These are converging lines of evidence, not interchangeable labels. A temporal link is powerful when clearly documented, but a characteristic cognitive profile can be equally important in a gradual subcortical presentation. Focal findings strengthen vascular attribution, while imaging establishes the burden and location of parenchymal injury. The clinician then asks whether the whole pattern is internally consistent. This is a causal formulation, not a checklist detached from the patient.

> **PITFALL**
>
> Do not equate a vascular risk factor or an incidental vascular lesion with vascular dementia. Risk, lesion and cause are different levels of inference. The cognitive syndrome and the cerebrovascular disease must fit each other.

9.2 Classification: several maps of the same territory

The classification varies with the purpose of the source. A clinical manual emphasises course and lesion distribution; an aetiological scheme separates infarction, small-vessel disease, hypoperfusion and haemorrhage; an examination-oriented table may preserve older vascular-dementia labels. These systems overlap substantially and should not be memorised as mutually exclusive diseases.

The contemporary clinical account describes **four** main forms: **poststroke neurocognitive disorder**, which becomes manifest immediately after a stroke; **subcortical ischaemic vascular neurocognitive disorder**; **multi-infarct cortical neurocognitive disorder**; and cortical-subcortical vascular neurocognitive disorder. The point of this division is practical. It asks whether the syndrome follows a recognised stroke, reflects predominantly deep small-vessel injury, accumulates with cortical infarction, or combines cortical and subcortical disease.

CLASSIFICATION LENS	NAMED FORMS	CLINICAL USE
Temporal onset	Acute-onset vascular dementia; poststroke neurocognitive disorder	Highlights a close relationship between a cerebrovascular event and cognitive change
Accumulating cortical injury	Multi-infarct dementia; cortical vascular neurocognitive disorder	Links progressive symptoms to repeated transient ischaemic episodes and accumulating infarcts
Deep small-vessel injury	Subcortical vascular dementia	Emphasises white-matter injury with relative preservation of the cerebral cortex
Combined distribution	Mixed cortical and subcortical vascular dementia	Recognises disease involving both compartments
Critical lesion	**Strategic infarct dementia**	Names the aetiological subtype without replacing lesion-specific assessment
Reduced perfusion	**Hypoperfusion dementia**	Classifies by the vascular mechanism rather than by cortical location alone
Bleeding	**Haemorrhagic dementia**	Separates haemorrhagic from ischaemic vascular injury
Residual categories	Other vascular dementia; vascular dementia unspecified	Allows classification when a more specific subtype is unsupported

Because these labels classify different axes, they need not behave as rival boxes. A label may describe timing, another the anatomical distribution and another the vascular mechanism. In a case formulation, state which axis each descriptor answers and select the term that best communicates the dominant clinical pattern. This avoids false precision when the available history and imaging support a broad vascular attribution but do not justify a narrowly named subtype.

A broader aetiological classification also names Alzheimer disease with cerebrovascular disease and hereditary vascular dementia, including CADASIL as a subtype label. Those labels matter because vascular and degenerative pathology may coexist, and because some vascular disorders are inherited. Their detailed diagnosis and management are not inferred from the label alone. Mixed dementia management is addressed with vascular dementia management elsewhere in this chapter, while the inherited and post-stroke vascular syndromes are covered in the Stroke, TBI, delirium and focal brain syndromes notes.

9.3 Clinical phenotypes: course, distribution and cognitive pattern

Classically, vascular dementia evokes abrupt onset, fluctuations and stepwise deterioration. Those features remain useful clues, especially when deficits change in relation to recognised cerebrovascular events. They are not universal. Acute-onset vascular dementia may follow several strokes or a severe stroke, while multi-infarct dementia produces progressive symptoms as

transient ischaemic episodes and infarcts accumulate. In contrast, small-vessel disease may produce gradual onset and slow progression. A good history therefore reconstructs change over time rather than forcing every patient into the stepwise template.

The cognitive phenotype is often uneven. Information processing, complex attention and frontal-executive functions may be prominent, so relatives may report slowness, difficulty shifting set, reduced initiation, loss of efficiency in multistep tasks and a need for repeated prompting. These complaints can be more revealing than an isolated report of forgetfulness. The examination should relate the pattern of cognitive inefficiency to focal neurological findings and to evidence of cerebrovascular disease, while remaining alert to an alternative or additional disorder.

Useful clinical clues can be organised without pretending that an individual clue is diagnostic:

- a cognitive change that is temporally linked to a cerebrovascular event;
- an uneven profile with prominent processing-speed, attention or executive impairment;
- focal neurological signs or symptoms that support brain vascular injury;
- a course whose fluctuations or increments correspond to vascular events rather than ordinary day-to-day variability.

The final clause is important. Fluctuation is a descriptive feature, not a vascular signature by itself. Similarly, somatic complaints and nocturnal confusion appear among the features sampled by the ischaemic scale, but neither establishes vascular causation. The pattern gains meaning only when it converges with the history, examination and imaging.

GOOD TO REMEMBER

Ask the informant to anchor each loss of function to an event, admission or new neurological deficit. A timeline often distinguishes a true event-linked decline from a vague retrospective impression that the course was “stepwise.”

9.4 The subcortical syndrome

Subcortical vascular dementia deserves separate attention because it is a common source of diagnostic error. It is associated with a history of hypertension and foci of white-matter injury, with the cerebral cortex relatively preserved. The resulting syndrome is not simply “memory loss with vascular risk.” It is a disorder of cognitive control, speed and retrieval, accompanied by characteristic behavioural and neurological features.

The described profile includes executive dysfunction, poor verbal fluency, psychomotor slowness, perseveration and inattention. Memory impairment is characterised by poor retrieval with recognition remaining intact. This distinction is clinically useful: the patient may struggle to generate stored information without help yet recognise it when the correct material is presented. That pattern points toward inefficient access and organisation rather than proving that information was never encoded. It supports a subcortical formulation but does not replace assessment of the entire syndrome.

Apathy and depression may accompany the cognitive changes. Pseudobulbar palsy may provide an additional neurological clue. The behavioural presentation can be misread as lack of effort because the patient initiates slowly, speaks little and needs prompting. Equally, psychomotor slowness can depress performance across many tasks, making the cognitive deficit appear globally worse unless time and executive demand are considered.

MNEMONIC

Slow, Set, Search, Spirit, Speech

The subcortical spine is **slow** processing and movement, impaired executive **set**, inefficient memory **search** with relatively intact recognition, reduced **spirit** through apathy or depression, and impaired **speech** output through poor verbal fluency, with pseudobulbar features as a neurological clue.

The diagnostic value lies in the configuration. Executive dysfunction plus slowed processing, impaired fluency and retrieval difficulty is more coherent than any isolated deficit. Neuroimaging then tests whether the distribution and burden of vascular injury can plausibly explain that configuration. The physics and general interpretation of imaging modalities are covered in the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

9.5 Contemporary diagnostic criteria

The **DSM-5-TR criteria for major or mild vascular neurocognitive disorder** begin by requiring the corresponding major or mild neurocognitive syndrome. Vascular aetiology is then supported through either a temporal relationship between cognitive deficits and cerebrovascular events, or a decline prominent in complex attention, including processing speed, and frontal-executive function. There must also be evidence of cerebrovascular disease from the history, physical examination or neuroimaging that is sufficient to account for the deficits. Another brain disease or systemic disorder must not explain the presentation better.

DIAGNOSTIC LAYER	QUESTION TO ASK	MEANING
Neurocognitive syndrome	Are criteria for major or mild neurocognitive disorder met?	Establishes that there is a syndrome to explain
Vascular clinical pattern	Is onset event-linked, or is decline prominent in complex attention, processing speed and frontal-executive function?	Provides a clinically coherent vascular phenotype
Cerebrovascular evidence	Does history, examination or neuroimaging show disease sufficient to account for the deficits?	Moves from association toward attribution
Exclusion	Is another brain disease or systemic disorder a better explanation?	Prevents premature closure
Certainty	Is the evidentiary threshold for probable disorder met?	Separates a supported vascular attribution from a possible one

Probable vascular neurocognitive disorder is diagnosed when there is neuroimaging evidence of significant parenchymal injury attributable to cerebrovascular disease, a documented temporal relationship between the syndrome and cerebrovascular events, or clinical and genetic evidence of cerebrovascular disease. If these supports are absent but the clinical criteria otherwise point to vascular aetiology, the diagnosis is **possible vascular neurocognitive disorder**. "Probable" therefore reflects the strength of causal evidence; it is not a synonym for greater cognitive severity.

MNEMONIC

TIME

Temporal event link or characteristic cognitive pattern; **I**njury from cerebrovascular disease sufficient to explain the deficit; **M**ajor or mild neurocognitive syndrome established; **E**xclude a better brain or systemic explanation.

9.6 Research and older examination criteria

The **NINDS-AIREN criteria** are the most widely used research criteria for vascular dementia. They require evidence of cerebrovascular disease from both clinical and neuroimaging findings, together with a temporal relationship between the vascular changes and the development of dementia. This makes them useful when a study needs a relatively well-defined vascular group. Across several clinical criteria sets, probable and possible vascular dementia diagnoses showed high specificity but low sensitivity. A strict research definition may therefore identify convincing cases while failing to capture every clinically relevant presentation.

Older manuals remain useful for understanding examination language. Under DSM-IV(R), dementia required either focal neurological signs and symptoms or evidence of cerebrovascular disease judged to be aetiologically relevant. The ICD-10 vascular dementia formulation required evidence of focal brain damage, aetiologically relevant cerebrovascular disease and an uneven distribution of cognitive deficits. The contrast is not cosmetic: the latter makes unevenness an explicit requirement, whereas the former allows diagnosis through either focal neurological features or relevant cerebrovascular evidence.

For current classification, the primary manual should be consulted for the exact wording. The supplied evidence supports the contemporary DSM cross-map and the older ICD comparison, but not a faithful reproduction of the current criteria. That limitation is preferable to silently importing remembered wording into a diagnostic standard.

TRAP **Do not present research criteria as though they were a sensitive bedside screen.** Their demand for converging clinical, imaging and temporal evidence improves specificity but can miss clinically plausible vascular presentations.

9.7 The ischaemic scale: what it can and cannot say

The **Hachinski Ischaemic Scale** was among the earliest criteria used in vascular-dementia research. It scores a vascular-looking clinical pattern through features such as abrupt onset,

stepwise deterioration, fluctuation, nocturnal confusion, somatic complaints, and vascular risks or signs. This list is illustrative rather than a complete reproduction of the weighted scale. The instrument has not been well validated, although interrater reliability can be good.

Its central limitation is conceptual. The scale estimates whether clinically significant cerebrovascular disease is present; it cannot demonstrate that the vascular disease caused the dementia. This distinction matters because multiple small infarcts frequently coexist with a primary degenerative dementia. A high vascular signal may therefore identify comorbidity rather than a pure vascular aetiology, which limits the scale's clinical usefulness.

A source reports that a score **greater than 4** was used as an exclusion criterion in a drug trial for Alzheimer disease because it was considered suggestive of vascular dementia. That is an applied trial threshold, not a validated general interpretation scheme for routine practice. It must not be converted into the familiar remembered low, mixed and high bands when the source does not supply those bands. Nor should the illustrative feature list be typeset as though it were the complete weighted item set.

The correct use of the scale is modest. It can prompt a careful vascular history and neurological examination, and it can summarise a cluster of vascular clues in a reproducible way. It cannot replace neuroimaging, establish lesion-cognition concordance, exclude coexisting degeneration or settle aetiology by arithmetic. Modern diagnostic reasoning therefore places the score beneath, not above, the clinical formulation.

PITFALL

A numerical score can create false certainty. Before acting on it, state which findings produced the score, whether cerebrovascular disease is objectively demonstrated, and whether that disease actually explains the cognitive phenotype.

9.8 A compact clinical and examination synthesis

A strong answer begins with the diagnostic principle, then classifies, phenotypes and applies criteria. Start by defining vascular dementia as a major neurocognitive syndrome attributable to sufficient cerebrovascular disease. Name the principal poststroke, subcortical, multi-infarct cortical and cortical-subcortical forms. Add the aetiological labels of strategic infarction, hypoperfusion and haemorrhage only to show the breadth of vascular mechanisms, not to launch into lesion syndromes owned elsewhere.

Next, describe the course without overgeneralising. Abrupt onset, stepwise deterioration and fluctuation are useful clues, but gradual progression occurs in small-vessel disease. Then give the characteristic cognitive profile: slowed information processing, impaired complex attention and frontal-executive function. For the subcortical form, add poor verbal fluency, psychomotor slowness, perseveration, inattention, retrieval-based memory impairment with intact recognition, apathy, depression and pseudobulbar palsy.

Finally, show how criteria convert description into diagnosis. Establish major or mild neurocognitive disorder; demonstrate either an event-linked onset or a characteristic attention-

executive decline; show cerebrovascular disease sufficient to account for the deficits; and exclude a better explanation. Distinguish probable from possible by the strength of imaging, temporal, clinical and genetic support. Research criteria demand clinical and imaging evidence plus temporal relatedness and may trade sensitivity for specificity. The ischaemic scale is an adjunct that detects vascular features, not a causal test.

> Memorise the hierarchy: cognitive syndrome, vascular phenotype, sufficient cerebrovascular evidence, temporal or anatomical concordance, exclusion of a better cause, then certainty level.

10 · Vascular dementia management, secondary prevention and mixed dementia

The marks lie in separating three questions that look deceptively similar. First, how should a person with vascular dementia receive ordinary dementia care? Second, which measures prevent further cerebrovascular injury? Third, when degenerative and vascular processes coexist, which component should drive treatment? A strong answer does not promise that control of vascular risk reverses established cognitive impairment. It explains that prevention of another vascular event is valuable in its own right, may prevent additional cognitive injury, and must be distinguished from a drug having a direct disease-modifying effect on vascular dementia.

Management is therefore broader than prescribing a cognitive enhancer. It begins with the care setting and immediate clinical problem, proceeds through secondary prevention, and ends with an integrated formulation of cognition, function, behaviour, vascular recurrence risk and support needs. The available evidence also demands restraint. **Treatment uncertainty in pure vascular dementia** is not therapeutic nihilism: management options are limited and chiefly address underlying cerebrovascular risk, while no available drug is formally licensed for vascular dementia in the United Kingdom. Mixed pathology changes the reasoning because a coexisting degenerative component may supply a separate indication for cognitive-enhancer treatment.

LOCUS	FOCUS	MODUS	REVIEW
CHOOSE THE CARE SETTING	DEFINE THE PHASE TARGET	COMBINE TREATMENT MODES	REASSESS BENEFIT AND HARM

10.1 The therapeutic frame: LOCUS, FOCUS and MODUS

LOCUS asks where care should occur. Stable cognitive impairment, vascular-risk review and longitudinal support usually belong in outpatient care. A sudden focal neurological change, abrupt loss of function or an acute safety problem requires an urgent medical pathway rather than routine adjustment of dementia medication. Inpatient care is justified by the severity and immediacy of the medical or behavioural problem, not by the diagnostic label alone. This prevents a common category error: treating a new clinical event as if it were merely another step in chronic cognitive decline.

FOCUS asks what is being treated in the present phase. The immediate focus is recognition of a new vascular or medical event and protection from harm. The short-term focus is medication reconciliation, clarification of vascular indications and restoration of a workable routine. The

medium-term focus is prevention of recurrent cerebrovascular injury while tracking cognition, function and behaviour. The long-term focus is sustained risk-factor care, realistic goals, caregiver partnership and review of whether each intervention still offers net benefit.

MODUS asks whether all relevant modes have been considered. Pharmacological care includes treatment of vascular risk and, in selected mixed cases, treatment directed at the degenerative component. Procedural care may be relevant for suitable carotid disease. Psychological care includes explanation, adaptation to deficits and support for adherence. Social care includes supervision, environmental structure and support for those carrying day-to-day responsibility. Detailed cognitive-enhancer pharmacology and general behavioural and psychological symptom management are taught elsewhere in this chapter; here the task is to decide whether vascular or mixed pathology changes their indication.

MANAGEMENT AXIS	QUESTION AT REVIEW	PRACTICAL OUTPUT
LOCUS	Is this stable follow-up or an acute change?	Outpatient review or urgent medical assessment
FOCUS, immediate	Is there new neurological or medical injury?	Identify and manage the acute problem
FOCUS, continuing	What can cause further vascular injury?	Individual secondary-prevention plan
MODUS, drug	Does each medicine have a defined indication?	Prescribe for stroke prevention, cognition or behaviour explicitly
MODUS, procedural	Is a vascular lesion amenable to intervention?	Appropriate specialist assessment
MODUS, psychosocial	What impairs adherence, function or safety?	Education, structure and caregiver support

RULE **Name the indication for every intervention.** A statin may be prescribed for vascular prevention, while a cognitive enhancer may be considered because of a degenerative component; neither should be described vaguely as a treatment that reverses vascular dementia.

10.2 Secondary prevention: prevent additional vascular injury

The core management task is **vascular secondary prevention**. In a person with established stroke, it is reasonable to address risk factors when the individual risk-benefit balance is favourable. The treatment plan should consider hypertension, hyperlipidaemia, smoking and diabetes; antiplatelet treatment when indicated by the history of stroke disease; and anticoagulation in atrial fibrillation. Significant carotid stenosis may warrant surgical assessment, with the cited reference identifying possible benefit when stenosis is **greater than 75%**. These measures are not a generic dementia prescription. Each follows from the person's vascular mechanism, history, contraindications and capacity to participate in monitoring.

The reasoning is causal but modest. A further stroke may add new cognitive, behavioural or functional disability to an already vulnerable brain. Reducing recurrent stroke risk can therefore protect the remaining level of function and may slow deterioration caused by repeated injury.

This does not prove that the same treatment removes established vascular lesions or directly alters the underlying dementia process. The candidate who preserves that distinction can reconcile apparently conflicting textbook statements without pretending that the evidence is cleaner than it is.

The **stroke-prevention versus dementia-modification distinction** is particularly important for statins and aspirin. Stroke literature supports treating appropriate vascular indications, yet the cited dementia evidence does not conclusively show that treatment of hyperlipidaemia with statins or clotting abnormalities with aspirin changes the incidence or progression of vascular dementia. A review cited in the same source found no studies supporting statins as a treatment for vascular dementia itself. Thus, prescribe a statin because hyperlipidaemia or cerebrovascular disease warrants it, not because the medicine has been established as a cognitive enhancer.

- Confirm the vascular indication and the event being prevented.
- Review bleeding risk, interactions, adherence and the burden of monitoring.
- Treat smoking and diabetes as part of the same vascular formulation.
- Use specialist assessment when carotid intervention is potentially relevant.
- At follow-up, reassess both recurrent events and everyday cognitive or functional change.

GOOD TO REMEMBER

A family may hear “there is no clearly effective drug for vascular dementia” as “nothing can be done.” Correct that misunderstanding. Prevention of another stroke, rational treatment of comorbidity, functional support and review of mixed pathology remain active treatment, even when restoration of lost cognition is not promised.

10.3 Cognitive enhancers: small signals, uncertain clinical importance

NICE 2006 guidance concluded that acetylcholinesterase inhibitors and memantine produced small cognitive benefits of uncertain clinical significance in mild-to-moderate vascular dementia, with insufficient evidence to support widespread use. This remains a useful examination formulation because it separates a detectable change on a cognitive outcome from a change that patients and caregivers can recognise in daily life. A statistically favourable result does not automatically establish meaningful functional benefit, global improvement or a routine indication.

A later evidence summary gives the same caution greater pharmacological detail. In vascular cognitive impairment, a systematic review found a slight cognitive benefit with donepezil and galantamine at the reviewed regimens. The vascular cognitive impairment review examined the standard Alzheimer regimens of **donepezil 5 mg or 10 mg and galantamine 16-24 mg**, tested here off their licensed indication. The size of change was considered unlikely to be clinically important. The higher reviewed donepezil regimen and the galantamine regimen were probably associated with more adverse events than placebo, while evidence for rivastigmine was less certain. These are trial regimens for vascular cognitive impairment, not a licence to reproduce the full dementia prescribing schedule or to bypass ordinary tolerability assessment.

The best synthesis is that **no drug is clearly effective in pure vascular dementia**, whereas acetylcholinesterase inhibitors may be beneficial in mixed dementia. That conclusion is compatible with selected trial signals if "effective" means a clinically worthwhile and broadly applicable treatment rather than any statistically detectable cognitive change. Before prescribing, ask what diagnosis supplies the indication, what outcome would count as success, how adverse effects will be recognised and when treatment will be reconsidered.

TRAP **Do not convert statistical efficacy into settled clinical effectiveness.** The cited analyses are not neatly concordant: one describes small cognitive benefits of uncertain clinical importance and insufficient support for widespread use, while another reports significant efficacy on cognitive outcomes and a favourable safety account. A further review reports more adverse events with some regimens. State the conflict, distinguish statistical from clinical significance, and do not resolve it by fiat.

PITFALL

Starting a cognitive enhancer simply because imaging shows vascular lesions confuses pathology with indication. First decide whether the clinical formulation is predominantly vascular, predominantly degenerative, or genuinely mixed. Then document the target outcome and review it rather than continuing treatment indefinitely by inertia.

10.4 Aggression and other behavioural presentations

Behavioural disturbance requires the same disciplined formulation. Clarify the behaviour, its setting, precipitating factors, risk, caregiver response and possible medical contributors before selecting medication. Environmental modification, communication strategy, routine and caregiver support remain foundational. The detailed assessment of behavioural and psychological symptoms, antipsychotic risk and general antipsychotic pharmacology belong respectively to the relevant dementia section and to the Schizophrenia: management, antipsychotics and adverse effects notes.

Evidence for vascular-specific pharmacological alternatives is sparse. A small case series reported that **gabapentin for aggression** was associated with reduced aggression in vascular dementia or mixed vascular and Alzheimer pathology at **200-600 mg/day**; some participants were subsequently able to stop antipsychotics. This was not a randomised trial, so it should be presented as a weak signal rather than a standard treatment algorithm. The indication attached to the cited regimen is aggression in vascular or mixed dementia, not uncomplicated cognitive decline.

In practice, weak evidence demands stronger review discipline. Define the target behaviour in observable terms, look for reversible precipitants, involve caregivers in monitoring, and stop or change an intervention when benefit is absent or harm outweighs it. Avoid presenting a small uncontrolled observation as equal to replicated comparative evidence. The appropriate answer is neither "never consider" nor "recommended treatment," but "limited evidence, individualised use, explicit target and close reassessment."

10.5 Mixed dementia: recognising more than one active process

Mixed dementia is a clinical formulation, not a compromise label. **Mixed dementia** commonly refers to the coexistence of vascular dementia and Alzheimer disease. The point is not merely that vascular lesions and degenerative pathology can both be found. The clinician must decide whether both plausibly contribute to the person's neurocognitive, functional or behavioural syndrome. A vascular history without clinical relevance should not be promoted automatically to a causal diagnosis, and a progressive amnestic syndrome should not make vascular contributions invisible.

Mixed pathology becomes especially important in later-onset presentations. **Late-onset mixed pathology** is more likely to be clinically recognised, and community samples often show more than one pathological contribution, sometimes alongside other diagnoses. Autopsy evidence makes the same caution concrete. In a European autopsy series, Alzheimer plus vascular pathology was reported in **5.5% overall, 10.6% above 90 years, and 5.2% at 60-69 years**. These are autopsy-series figures, not the prevalence of a clinical diagnosis in an outpatient service. They support diagnostic openness, not indiscriminate use of the mixed label.

Recognition rests on convergence. The temporal course may contain both progressive decline and changes plausibly linked to cerebrovascular disease. The cognitive profile may combine cortical impairments with slowed processing, attentional difficulty or executive dysfunction. Medical history and neuroimaging may demonstrate cerebrovascular disease, but their contribution must be interpreted in relation to symptoms and function. Detailed vascular-dementia subtypes, diagnostic criteria and the ischaemic score are covered in the preceding vascular-dementia section; imaging techniques themselves belong to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

MNEMONIC

Two tracks, one patient

Keep both tracks visible: **prevent further vascular injury** and **treat the predominant dementia process**. The tracks interact, but their indications and expected outcomes should remain explicit.

10.6 Diagnostic formulation across the major classifications

The international diagnostic classification recognises a dedicated **ICD-11 mixed Alzheimer and vascular category**, coded **6D80.2**. It requires that the diagnostic requirements for dementia due to Alzheimer disease are met and that part of the neurocognitive, functional or behavioural syndrome is attributable to coexisting cerebrovascular disease demonstrated by clinical history, medical testing or neuroimaging. The expected course remains progressive, with possible combined impairment in cortical functions and in attention, processing speed or executive and frontal functions.

The **DSM-5-TR comorbidity formulation** reaches the same clinical destination by a different route. When major or mild neurocognitive disorder due to Alzheimer disease co-occurs with major or mild vascular neurocognitive disorder, both diagnoses should be made. For examination

purposes, the distinction is straightforward: the international classification supplies a named mixed category, while the American manual records the co-occurring aetiological diagnoses. Neither approach permits the clinician to infer causal contribution merely from an incidental lesion.

QUESTION	EVIDENCE SOUGHT	INTERPRETIVE TASK
Is a degenerative dementia syndrome present?	Progressive cognitive and functional history	Establish the degenerative component rather than assuming it
Is cerebrovascular disease present?	Clinical history, medical tests and neuroimaging	Document vascular disease
Does vascular disease contribute?	Temporal, anatomical and cognitive concordance	Link evidence to symptoms and function
How should it be coded?	Chosen diagnostic system	Use the mixed category or record both aetiological diagnoses
What drives treatment?	Predominant cause and active risks	Match each intervention to its indication

10.7 Integrated management when pathology is mixed

The governing principle is **management according to the predominant cause**. "Predominant" does not mean "only." It identifies which process best explains the current syndrome and therefore which disease-specific treatment framework should lead. A predominantly vascular picture places prevention of further cerebrovascular injury at the centre. A substantial degenerative component may justify cognitive-enhancer treatment under that indication. In either case, active vascular risks still deserve attention, and cognitive, behavioural, functional and caregiver outcomes still require longitudinal review.

This formulation should be written as a short causal map. State the evidence for a progressive degenerative process, the evidence for cerebrovascular disease, and why each is judged clinically contributory. Then attach each intervention to one of those components. Such documentation prevents two opposite errors: withholding potentially useful treatment because vascular disease is present, and prescribing as though a mixed label proves that every dementia medicine will help.

A practical mixed-dementia plan therefore combines prevention, selective symptomatic treatment and supportive care. It also preserves proportionality. The potential benefit of vascular prevention depends on the relevant vascular indication and risk-benefit profile. The potential benefit of a cognitive enhancer depends on the degenerative contribution, modest expected effect and tolerability. Psychological and social interventions should address function, communication, adherence and caregiver strain without waiting for perfect aetiological certainty. Goals should remain observable and revisable.

Follow-up should return to the causal map rather than merely renew prescriptions. Ask whether any new vascular event has occurred, whether cognition or everyday function has changed, whether the target behaviour remains present, and whether the family can still carry out the plan.

Review adherence and adverse effects alongside the burden of appointments and monitoring. If the balance of disability changes, reconsider which process is predominant and whether the present treatment priorities still fit. This longitudinal method is especially important when the course contains both gradual decline and superimposed vascular injury, because a static label can conceal a changing clinical problem.

> In vascular and mixed dementia, prevent further vascular injury, identify the predominant cause, and never mistake a small cognitive signal for proven disease modification.

10.8 How to write the management answer

Begin with setting and phase: state whether the patient is stable or acutely changed, then use LOCUS, FOCUS and MODUS to organise care. Next, describe vascular secondary prevention with individual risk-benefit assessment. Explicitly separate reduction of recurrent stroke from direct treatment of established dementia. Then discuss cognitive enhancers cautiously: small cognitive benefits may occur, clinical importance is uncertain, evidence summaries conflict, and routine widespread use in pure vascular dementia is not established.

Finish by defining mixed dementia, explaining how vascular and degenerative evidence are linked to the clinical syndrome, and stating the diagnostic-system formulation. Management follows the predominant cause while continuing to address both active processes. If behavioural symptoms are mentioned, assess them before prescribing and grade the evidence honestly. This sequence earns marks because it is clinically reasoned rather than a list of drugs, and because it makes uncertainty visible without turning uncertainty into inaction.

11 · Frontotemporal dementia: the behavioural variant and the primary progressive aphasias

The diagnostic centre is the leading network failure, not memory loss by default.

Frontotemporal degeneration can announce itself through altered social conduct and executive control, or through a progressive language-led syndrome. The high-yield task is therefore to identify what changed earliest, describe the dominant phenotype precisely, and then apply the relevant diagnostic frame without forcing every presentation into an amnestic dementia template.

The behavioural variant and the primary progressive aphasias belong in a continuous clinical argument because behaviour, social cognition, executive control and language can overlap as degeneration spreads. They must nevertheless remain distinguishable at presentation. A vague label such as “frontal dementia” loses marks and clinical information. A strong formulation names the syndrome, shows the defining findings, states the relatively preserved functions, and acknowledges when the language phenotype points away from frontotemporal pathology.

2-31 per 100,000 POPULATION PREVALENCE ESTIMATES	**20-25%** PRESENT AFTER AGE 65	**ABOUT 5%** UNSELECTED DEMENTIA AUTOPSY SERIES	**APPROXIMATELY 60%** BEHAVIOURAL-VARIANT PRESENTATION

11.1 The frontotemporal clinical spectrum

Frontotemporal neurocognitive disorder is uncommon in population terms but disproportionately important in psychiatric practice because its earliest manifestations may look like a disorder of motivation, impulse control, personality or communication. The behavioural variant is the commonest presentation, accounting for **approximately 60%** of cases in the cited diagnostic-manual account. A meaningful minority present after 65 years, so older age must not be used to dismiss the diagnosis.

The accompanying figure should be read as a phenotype map. At the behavioural pole, early dysfunction concerns social comportment, empathy, motivation, repetitive behaviour, eating and executive control. At the language pole, the dominant impairment may concern word meaning, grammatical or motor production, or retrieval with phrase and sentence repetition. These are clinical entry points rather than sealed compartments. Behavioural change may appear during a language-led illness, and language difficulty may emerge during a behaviour-led illness.

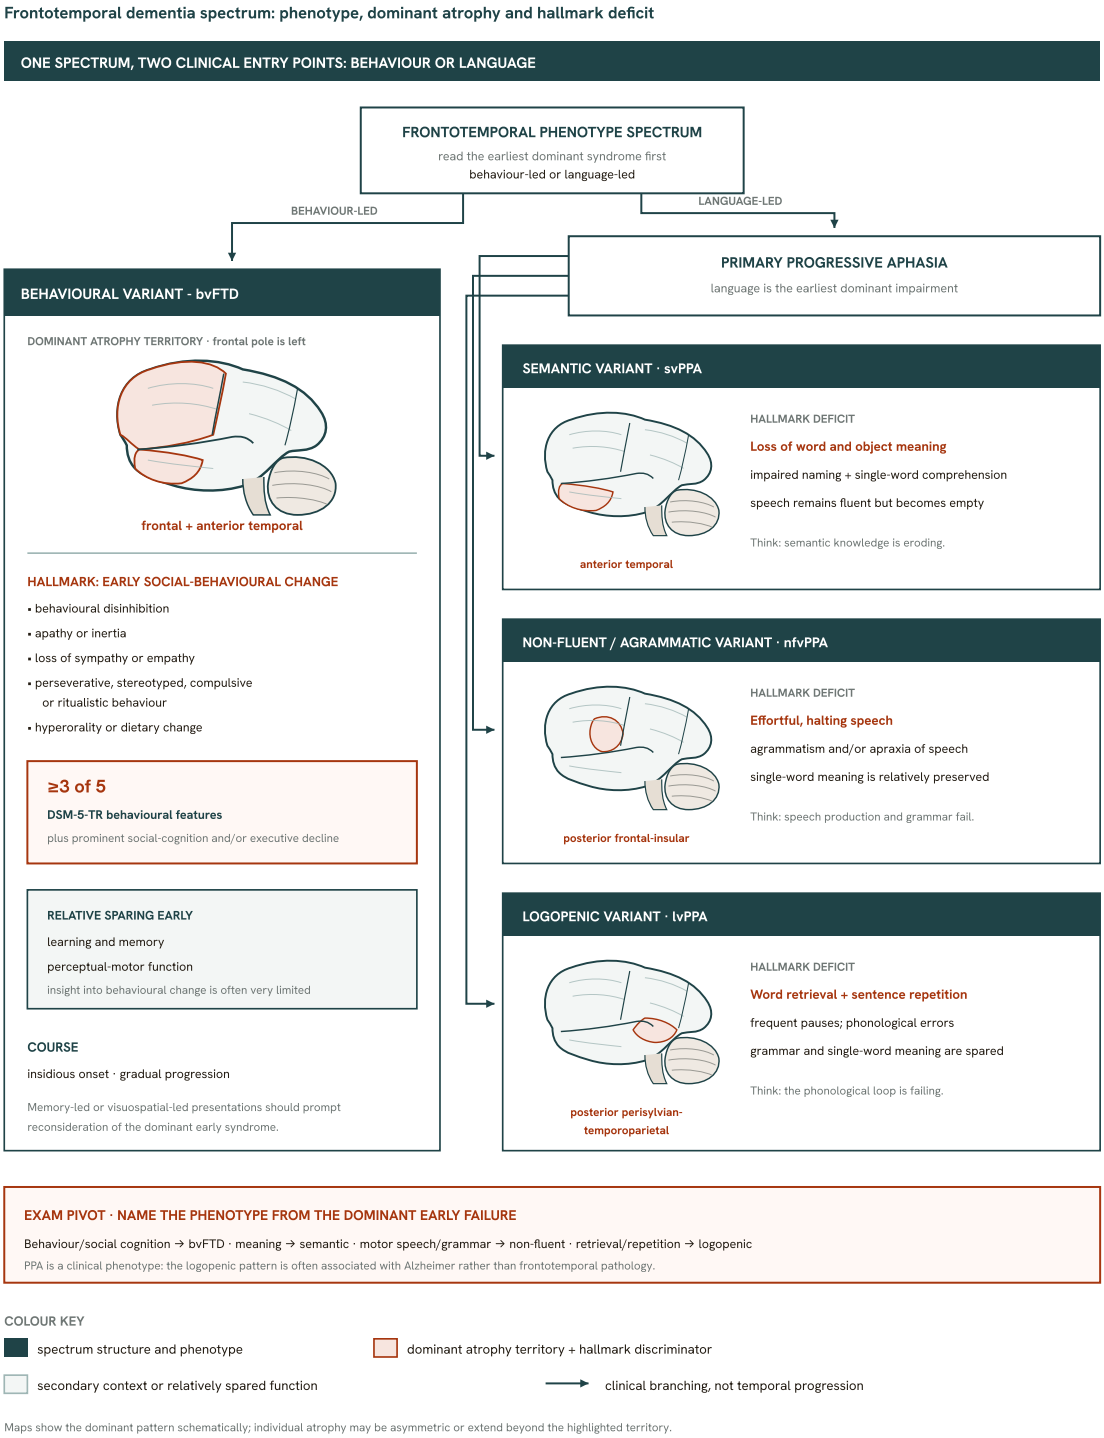


Fig 32.5 **The frontotemporal dementia phenotype map.** Behaviour-led disease maps to frontal and anterior temporal involvement; language-led disease separates into semantic, non-fluent/agrammatic and logopenic primary progressive aphasia by the dominant early language failure and its characteristic network. The logopenic phenotype is often associated with Alzheimer pathology rather than frontotemporal lobar degeneration.

Course reinforces the seriousness of early recognition. Median survival is **6-11 years after symptom onset** in the cited account, with faster decline and shorter survival than in typical Alzheimer disease. That figure describes a group trajectory, not an individual prognosis. Clinically, the useful implication is that apparently “social” or “personality” symptoms can

represent a progressive neurological syndrome with major consequences for judgement, relationships, occupation and care planning.

RULE **Classify the earliest dominant syndrome before interpreting later overlap.** Early social-executive change supports behavioural-variant frontotemporal dementia; progressive language-led change requires analysis of meaning, production, retrieval and repetition before a subtype is assigned.

11.2 Behavioural variant: recognise change from the person's baseline

The **behavioural variant of frontotemporal dementia** is defined clinically by a progressive alteration in how the person initiates action, regulates conduct, understands other minds and adapts behaviour to context. The word "change" matters. Lifelong bluntness, unconventionality or rigid habits are not equivalent to a newly progressive syndrome. The history should reconstruct premorbid personality and compare it with current behaviour across home, work and social settings.

Apathy or inertia may be described as loss of spontaneous activity, reduced initiation, diminished interest in former responsibilities or prolonged passivity unless another person prompts action. The patient may not complain of distress about this loss. Relatives often report that the person "could do it but does not start," which points the interview toward initiation and goal-directed behaviour rather than ability alone. The examination should test whether prompting briefly restores action and whether the person appreciates the consequences of disengagement.

Behavioural disinhibition ranges from loss of manners and tact to plainly socially inappropriate conduct. The key is failure to use social context to inhibit an impulse. Familiarity with strangers, intrusive comments, impulsive purchases or disregard for interpersonal boundaries may be more diagnostically revealing than a conventional memory complaint. The clinician should document the actual behaviour and its novelty rather than substitute a moral description such as "difficult" or "selfish."

Loss of sympathy or empathy is more than irritability. The person may fail to register another person's emotional state, respond coldly to suffering, or lose the reciprocal concern previously evident in close relationships. Family members experience this as a profound personality change. Because the patient may explain the behaviour fluently, the clinician must compare verbal explanation with conduct in real situations. Social knowledge stated in the abstract does not guarantee that it guides behaviour.

Perseverative, stereotyped or compulsive behaviour includes repeated acts, fixed routines and ritualised patterns that continue despite reduced usefulness. The form matters more than whether the act resembles an ordinary habit. Ask what happens if the routine is interrupted, whether the behaviour has become narrower or more frequent, and whether it serves an understandable anxiety-reducing purpose. In bvFTD, repetitive conduct can coexist with striking indifference to its impact.

Hyperorality and dietary change may appear as altered food preference, excessive oral exploration or a new pattern of eating that relatives identify as uncharacteristic. This feature is not merely "good appetite." Its diagnostic weight comes from novelty, repetition and its place within the wider behavioural-executive syndrome. Questions about food seeking, preference shifts, table behaviour and oral habits are often more productive than asking whether appetite has increased.

Limited insight is characteristic. Individuals may minimise change, rationalise serious consequences or regard relatives as unnecessarily restrictive. This does not make the interview futile; it changes the evidentiary balance. Collateral history becomes central, while the patient's account remains valuable for language, thought form, emotional response and awareness of disagreement. Executive testing may show impairment with relative sparing of memory and visuospatial functioning, but a neat test profile is not required for family-observed behavioural decline to deserve investigation.

GOOD TO REMEMBER

Ask the informant for concrete before-and-after examples. "He has changed" is weak evidence; a new loss of tact, initiation, empathy, behavioural flexibility or dietary restraint establishes the domain, context and functional consequence of change.

11.3 Diagnostic criteria and probability language

The **DSM-5-TR behavioural-variant criteria** begin only after criteria for major or mild neurocognitive disorder are met. They then require insidious onset and gradual progression. The behavioural syndrome requires **at least three of five behavioural features**, together with prominent decline in social cognition and/or executive abilities. The feature set is:

- behavioural disinhibition;
- apathy or inertia;
- loss of sympathy or empathy;
- perseverative, stereotyped, compulsive or ritualistic behaviour; and
- hyperorality and dietary change.

DIAGNOSTIC LAYER	REQUIRED CONTENT	WHY IT MATTERS
Entry syndrome	Major or mild neurocognitive disorder criteria are met	Behavioural symptoms alone do not establish the neurocognitive diagnosis
Course	Insidious onset and gradual progression	Frames the presentation as a progressive degenerative syndrome
Positive phenotype	At least three of the five specified behavioural features, plus prominent social-cognitive or executive decline	Combines observable behaviour with the affected cognitive systems
Relative preservation	Learning and memory, and perceptual-motor function, are relatively spared	Prevents reflex classification as a primarily amnestic or perceptual-motor syndrome
Exclusions	The picture is not better explained by cerebrovascular disease, another neurodegenerative disease, substance effects or another disorder	Requires an aetiologic differential rather than checklist diagnosis
Probability	Probable when a causative genetic mutation is demonstrated or neuroimaging shows disproportionate frontal or temporal involvement; otherwise possible	Separates a compatible clinical syndrome from one with specified aetiologic support

"Relative sparing" is frequently misread as "normal." A patient may have some forgetfulness, inefficient learning or visuospatial difficulty and still show a clearly disproportionate behavioural-executive syndrome. The task is comparative: determine which domain drove the illness, which impairment dominates current disability, and whether the overall pattern satisfies the stated framework. Later generalisation does not erase the presenting phenotype.

The International Behaviour-variant FTD Criteria Consortium revisions published in 2011 use a hierarchical classification of possible, probable and definite bvFTD, supported across the hierarchy by radiological, genetic and neuropathological information. That framework should not be silently merged with the manual formulation above. Its category names should be attributed to the consortium framework rather than mapped onto the manual's probable and possible split.

TRAP **Do not diagnose bvFTD by counting colourful behaviours alone.** The manual requires the neurocognitive syndrome, insidious progressive course, the threshold behavioural cluster, prominent social-cognitive or executive decline, relative sparing and appropriate exclusions. A behaviour checklist detached from course and cognition is incomplete.

11.4 Clinical examination, function and associated motor signs

Assessment is strongest when history, examination and everyday function converge. Begin with chronology: what was the earliest convincing change, who noticed it, and how has it progressed? Then sample initiation, inhibition, cognitive flexibility, rule maintenance, social judgement and

awareness of deficit. A calm, superficially cooperative interview can obscure severe dysfunction if the clinician does not ask how the person behaves outside the consulting room.

Collateral interviewing should explore work errors, financial judgement, vulnerability to exploitation, unsafe impulsivity, eating behaviour, interpersonal rupture, neglect of responsibilities and response to feedback. These examples are not separate diagnostic criteria. They are ways of showing that the core behavioural and executive changes have crossed into real function. When accounts conflict, document the discrepancy itself. Limited insight can be part of the syndrome, but informant frustration can also colour interpretation.

The bedside cognitive profile should not be reduced to a total screening score. Observe whether the patient grasps the interpersonal setting, modifies an answer after correction, sustains a rule, generates alternatives and recognises another person's perspective. Memory and perceptual-motor performance may be relatively spared in the formal diagnostic profile. General cognitive-screen administration and scoring belong to the Psychometry, Assessment and Descriptive Psychopathology notes; here their role is to describe the pattern and severity, not to replace clinical formulation.

Motor examination remains relevant even in an apparently behavioural presentation. About **20%** of individuals with bvFTD show signs of parkinsonism in the cited account. Motor neuron disease can also coexist and should be named when weakness, wasting or fasciculations accompany the cognitive-behavioural syndrome, but its frequency, molecular basis and detailed overlap belong with the dedicated frontotemporal molecular and motor-neuron-disease discussion. The practical point here is to examine movement, tone, gait, bulk and power rather than treating bvFTD as purely psychiatric.

> **PITFALL**
>
> Do not allow fluent conversation or adequate bedside recall to overrule a compelling history of progressive social-executive failure. Conversely, do not convert any late-life disinhibition into bvFTD without demonstrating progression, cognitive-system involvement and exclusion of a better explanation.

11.5 Primary progressive aphasia: classify the failed language operation

In progressive aphasia, "speech difficulty" is only the opening observation. The clinician must determine whether the central failure concerns semantic knowledge, grammatical or motor production, or retrieval with repetition. The **primary progressive aphasia** presentation has gradual onset and a progressive language-led course. Connected speech, confrontation naming, single-word comprehension, sentence comprehension, repetition, reading, writing, object knowledge, grammar, articulation and prosody should be sampled in a coordinated examination.

Taxonomy requires explicit attribution. The DSM-5-TR account describes **two commonly described language subtypes** within frontotemporal neurocognitive disorder: semantic and agrammatic or nonfluent. It separately names logopenic progressive aphasia, links it to left temporoparietal dysfunction and states that Alzheimer pathology often causes it. Other major

textbook conventions discuss semantic, nonfluent or agrammatic, and logopenic presentations together as the recognised PPA variants. The framing differs, while the clinical phenotypes remain useful.

The ICD-11 clinical description offers a compact discriminator: primary deficits concern word meaning in the semantic subtype, word production in the agrammatic subtype, and word finding in the logopenic subtype. This is an entry map, not a complete bedside diagnosis. Each label must be tested against associated deficits and relative preservations, because several patients can all pause for words while failing for different reasons.

AXIS	SEMANTIC VARIANT	NONFLUENT OR AGRAMMATIC VARIANT	LOGOPENIC VARIANT
Core failure	Confrontation naming and single-word comprehension	Agrammatism or effortful, halting speech	Single-word retrieval plus sentence and phrase repetition
Connected speech	Fluent and grammatically correct despite loss of semantic precision	Nonfluent, effortful; agrammatism is usual but not invariable	Slow because of word-finding pauses, without defining agrammatism
Relative preservation	Repetition and speech production are usually spared	Single-word comprehension and object knowledge are spared	Single-word comprehension, object knowledge and motor speech are spared; articulation and prosody are preserved
Associated clue	Impaired object knowledge, surface dyslexia or dysgraphia may occur	Apraxia of speech is common, and complex syntax may be difficult	Phonological errors may occur in speech and naming
Typical atrophy or dysfunction	Bilateral anterior temporal, usually greater on the left	Predominantly left inferior frontal and insular	Left temporoparietal dysfunction
Usual pathology framing	Associated with the FTD spectrum	Associated with the FTD spectrum	Associated with Alzheimer disease

MNEMONIC

Meaning - Making - Repeating

Meaning fails in semantic PPA, speech **making** fails in nonfluent or agrammatic PPA, and **repeating** phrases or sentences fails alongside word retrieval in logopenic PPA. The mnemonic directs examination; the full subtype still depends on the wider pattern.

11.6 Semantic variant PPA: words remain fluent but lose meaning

The **semantic variant of primary progressive aphasia**, also called semantic dementia, has impaired confrontation naming and impaired single-word comprehension as its core. Speech can remain fluent, well articulated and grammatically correct, yet become progressively empty or

imprecise because semantic knowledge is degraded. This contrast is crucial: fluency describes the mechanics and flow of speech, not the richness or accuracy of meaning.

At the bedside, naming failure should be tested alongside comprehension. A person may not merely fail to retrieve "apple"; they may substitute a broad category such as "food," and later may refer to nouns as "things." Ask the patient to identify, define and match words and objects rather than relying on picture naming alone. Impaired object knowledge supports loss of the underlying concept, whereas preserved repetition and speech production help separate semantic breakdown from a production disorder.

Surface dyslexia and dysgraphia may accompany the syndrome. They show that semantic knowledge contributes to reading and spelling patterns that cannot be managed reliably by a simple sound-based route. The finding should be interpreted within the language profile rather than treated as an isolated literacy deficit. Educational history, language proficiency and premorbid reading ability remain essential context in multilingual Indian practice, especially when testing is performed in a language that is not the patient's most familiar one.

The typical structural association is atrophy in the ventral and lateral portions of the bilateral anterior temporal lobes, usually greater on the left. Predominantly right-sided semantic degeneration can look different: emotional detachment, loss of empathy, compulsions, impaired emotional processing and difficulty recognising familiar faces may appear early. This is the bridge back to the behavioural spectrum. It also explains why a language-only interview can miss the syndrome when the right temporal presentation is socially dominant.

11.7 Nonfluent or agrammatic PPA: production becomes effortful

The **nonfluent or agrammatic variant of primary progressive aphasia**, also called progressive nonfluent aphasia, requires one of two core findings: agrammatism in language production, or effortful and halting speech. The patient may speak slowly and labour over utterances. Agrammatism appears as omission or misuse of grammatical structure, while impaired motor planning for speech can distort the route from intended utterance to articulated output.

Apraxia of speech is common in the detailed textbook description, but language and motor speech should still be examined separately. Ask whether grammatical organisation fails independently of motor speech and whether halting output reflects language formulation, speech production, or a mixture. Laboured output should not automatically be labelled agrammatism, and grammatical simplification should not automatically be explained by articulatory difficulty.

Comprehension of syntactically complex sentences may be impaired, even while single-word comprehension and object knowledge remain relatively spared. This dissociation is clinically elegant and frequently examined. A patient may understand the words individually yet fail to compute who did what to whom when grammatical structure carries the meaning. Word-finding difficulty is generally milder than in semantic PPA, although mild anomia may occur and may be more noticeable for verbs.

The associated atrophy is predominantly left inferior frontal and insular. Behavioural change is usually limited early, and individuals may retain more insight and adherence to social norms than those with bvFTD. Behavioural symptoms may follow if degeneration extends into right frontal regions. Thus, preserved social conduct does not argue against a frontotemporal-spectrum language syndrome; it helps identify the current phenotype and stage of network involvement.

11.8 Logopenic PPA: retrieval and repetition expose the distinction

The **logopenic variant of primary progressive aphasia** has impaired single-word retrieval in spontaneous speech and naming, together with impaired repetition of sentences and phrases. Speech becomes slow because the patient searches for words, not because articulation is effortful or grammar is collapsing. Phonological errors may occur, but articulation, prosody and motor speech are relatively preserved.

Repetition is the decisive stress test. A patient with semantic PPA may repeat material despite not understanding key words; a patient with nonfluent or agrammatic PPA may struggle because production is effortful; a patient with logopenic PPA shows the characteristic conjunction of retrieval pauses and poor repetition of phrases and sentences while single-word comprehension and object knowledge remain relatively intact. The examiner should therefore describe the reason repetition fails, not merely record "impaired."

The major nosological caution is pathological. Logopenic PPA is associated with left temporoparietal dysfunction and is commonly linked to Alzheimer pathology, whereas semantic and nonfluent or agrammatic PPA are associated with the FTD spectrum. Calling every progressive aphasia "frontotemporal dementia" therefore confuses a clinical language phenotype with its usual underlying disease. Conversely, acknowledging the Alzheimer association should not erase the value of the logopenic clinical label.

TRAP **Do not present logopenic PPA as unambiguously the same kind of FTD-spectrum subtype as semantic and nonfluent PPA.** It belongs in the clinical PPA comparison, but the diagnostic manual and textbook pathology framing link it more often to Alzheimer disease.

A compact synthesis is to listen for fluency, then test meaning, grammar, motor speech, retrieval and repetition. Fluent but semantically empty speech points toward semantic PPA. Effortful, halting or grammatically impoverished production points toward nonfluent or agrammatic PPA. Word-finding pauses with disproportionately impaired phrase and sentence repetition point toward logopenic PPA. In every case, the subtype is a pattern of deficits and preservations, not a synonym for any isolated naming failure.

Memorise the organising contrast: bvFTD leads with progressive social-executive and behavioural change; semantic PPA loses word meaning, nonfluent or agrammatic PPA loses fluent grammatical production, and logopenic PPA combines word-retrieval difficulty with impaired phrase and sentence repetition and is associated with Alzheimer disease.

12 · Frontotemporal dementia: molecular pathology, genetics and the motor-neurone-disease overlap

The examination hinge is a three-way mapping. A frontotemporal clinical syndrome must be related to its dominant molecular pathology, its possible genetic cause, and the possibility of motor-neurone disease. These are related layers, but they are not interchangeable. A behavioural or language phenotype describes the patient; a proteinopathy describes the tissue; a pathogenic variant or repeat expansion describes a cause. The marks lie in moving accurately between those layers without claiming that phenotype alone proves genotype or pathology.

This framework also explains why frontotemporal dementia and amyotrophic lateral sclerosis cannot be kept in separate conceptual boxes. Some families and patients express predominantly cognitive-behavioural disease, some predominantly motor disease, and some both. Shared molecular pathology, particularly **transactive response DNA-binding protein pathology**, provides the bridge. The practical consequence is reciprocal vigilance: assess motor function in frontotemporal syndromes and assess behaviour, executive function and language in motor-neurone disease.

12.1 From clinical syndrome to molecular disease

Frontotemporal lobar degeneration is the pathological umbrella underlying several frontotemporal clinical phenotypes. Molecular classification asks which abnormal protein predominates in neuronal or glial inclusions. The broadest useful division is between **FTLD-tau**, a group with tau-positive inclusions, and the tau-negative group with ubiquitin and p62-positive inclusions, most often represented by **FTLD-TDP**. A smaller **FTLD-FUS** category is particularly relevant to behavioural presentations. These labels are pathological, not alternative names for the behavioural variant or a primary progressive aphasia.

RULE **Keep phenotype, protein and gene on separate lines of the answer.** State the clinical syndrome first, then the likely proteinopathy, then the associated genetic causes. A strong association supports a hypothesis; it does not make bedside phenotype a molecular assay.

The distribution is broad enough to make all of these categories important. Tau-positive disease accounts for **36-50%** of frontotemporal lobar degeneration, TDP-positive disease represents about **50%** of frontotemporal dementia, and FUS pathology accounts for about **10%** of behavioural-variant frontotemporal dementia. The denominators differ, so these figures should not be forced into a single pie chart. Their value is to show that tau and TDP pathology dominate, while FUS is a smaller but examinable group.

36-50%	ABOUT 50%	ABOUT 10%	30-60%
FTLD WITH TAU PATHOLOGY	FTD WITH TDP PATHOLOGY	BVFTD WITH FUS PATHOLOGY	FTD WITH SUBCLINICAL MOTOR EVIDENCE

Within the tau-positive group sit classic tauopathies including Pick disease, corticobasal degeneration, progressive supranuclear palsy and dementia caused by a tau-gene mutation. Within the tau-negative ubiquitin/p62-positive group, intracellular TDP-positive inclusions are

the common pattern. This first-pass taxonomy is more useful than memorising disconnected lists because it predicts where the major genes fit.

12.2 Gene-to-protein mapping

The common familial causes cluster around three major genes: **MAPT**, **GRN**, also called PGRN in some texts, and the **C9ORF72** repeat expansion. The first maps to tau pathology. The progranulin gene maps to TDP pathology despite not encoding TDP itself. The repeat expansion is likewise associated with TDP pathology but may also generate ubiquitin/p62-positive, TDP-negative inclusions, especially in the cerebellum. Thus, a gene may predict the inclusion class without being the protein deposited in the inclusion.

GENETIC CAUSE	DOMINANT PATHOLOGICAL ASSOCIATION	CLINICAL OR ANATOMICAL CLUE
Tau-gene mutation	Tau-positive inclusions	Behavioural presentation is common; onset tends to be younger, with symmetric anteromedial temporal atrophy
Progranulin-gene mutation	TDP-positive inclusions	Behavioural and neuropsychiatric features are common; language and extrapyramidal involvement may appear, with asymmetric frontotemporal-parietal atrophy
Major repeat expansion	TDP-positive plus ubiquitin/p62-positive, TDP-negative inclusions	Behavioural disease, motor-neurone disease, or both; more generalised atrophy with thalamic and sometimes cerebellar involvement
VCP mutation	TDP-positive inclusions	Behavioural change, executive dysfunction and aphasia; may occur with motor disease, **inclusion-body myopathy** or Paget disease of bone
TARDBP mutation	TDP-related disease	Behavioural or semantic presentation, with or without motor-neurone disease
CHMP2B mutation	Ubiquitin/p62-positive but TDP-negative inclusions	Behavioural frontotemporal dementia or combined frontotemporal and motor disease
TBK1 loss-of-function mutation	FTD-MND spectrum association	Motor disease may accompany semantic or non-fluent aphasic presentations
FUS-related disease	FUS proteinopathy	Combined frontotemporal and motor-neurone phenotype or amyotrophic lateral sclerosis

The table is a map of associations, not a diagnostic algorithm. A behavioural phenotype can arise from more than one genetic or protein category, and the same gene can produce different phenotypes. Conversely, apparently sporadic disease does not establish absence of a genetic cause. Family history, neurological examination, neuroimaging and appropriate genetic assessment must be interpreted together.

The map matters because each layer answers a different question. Phenotype guides what is examined and followed; the atrophy pattern tests anatomical coherence; pathology explains the protein class; genetics may identify an aetiological route and inform relatives. Collapsing these questions produces false certainty. Keeping them separate allows convergence: a result becomes persuasive when the history, examination, imaging, pedigree and molecular finding point in the same direction, while discordance becomes a reason to recheck assumptions rather than discard inconvenient evidence.

MNEMONIC

Three layers: Face, Filing, Fault

Face is the clinical phenotype seen at the bedside; **Filing** is the protein inclusion category used by pathology; **Fault** is the causal genetic lesion. The mnemonic prevents the common error of treating one layer as proof of another.

12.3 The major familial patterns

The tau-gene form tends to present with a behavioural syndrome and at a younger age. The progranulin form also commonly produces a behavioural syndrome, but neuropsychiatric symptoms, early language involvement and extrapyramidal features are notable. Around half of affected patients in the cited series have extrapyramidal symptoms. The repeat-expansion form commonly produces behavioural frontotemporal dementia, amyotrophic lateral sclerosis, or their combination; apathy, disinhibition, impaired empathy and executive dysfunction are characteristic behavioural features. These are tendencies, not mutually exclusive templates.

The numerical contribution of a gene depends on the denominator. Progranulin loss-of-function mutations may account for **up to 5-10% of all FTD**, while a separate familial-FTD table assigns this gene a larger share of familial cases. Those statements are compatible because “all FTD” and “familial FTD” are different populations. This is exactly why an answer should write the denominator beside the proportion rather than listing a naked percentage.

Use the phenotype to prioritise, not to close, the differential. A predominantly behavioural presentation can arise from the major tau, progranulin or repeat-expansion causes. Early language involvement with asymmetric frontotemporal-parietal atrophy supports the progranulin pathway, while a combined behavioural and motor-neurone phenotype makes the major expansion especially relevant. Behavioural change, executive dysfunction and aphasia accompanied by myopathy or bone disease form a more distinctive clue to the valosin-containing-protein pathway.

Less common associations broaden the phenotype. A transactive DNA-binding-protein gene mutation may produce behavioural or semantic disease with or without motor-neurone involvement. The charged multivesicular-body-protein gene is unusual because its inclusions are ubiquitin/p62-positive but TDP-negative. A loss-of-function kinase mutation can produce combined frontotemporal and motor disease, including semantic and non-fluent aphasic presentations associated with amyotrophic lateral sclerosis. These rarer patterns are best retained as discriminating clues after the major trio is secure.

12.4 The repeat expansion as the FTD-ALS bridge

One lesion unifies the continuum especially clearly. The major hexanucleotide repeat expansion is the most common genetic cause of both familial frontotemporal dementia and familial amyotrophic lateral sclerosis. Its usual clinical expressions are behavioural frontotemporal dementia, motor-neurone disease, or a combination. Its underlying pathology is characteristically TDP-related, although the additional cerebellar ubiquitin/p62-positive, TDP-negative inclusions remind the reader that its tissue signature is not completely uniform.

Psychosis is a useful clinical flag in this genetic context. Hallucinations and delusions occur in 30-50% of expansion carriers with frontotemporal disease in the cited account. This does not mean that psychosis identifies the mutation, nor that every late-onset psychotic presentation is neurodegenerative. It means that new psychotic symptoms should not be allowed to obscure a progressive behavioural-executive syndrome, especially when motor disease or a relevant family history is present. The formal differentiation of behavioural-variant frontotemporal dementia from primary psychiatric disorders belongs in the dedicated differential-diagnosis section.

Neuroimaging offers another clue. Compared with other genetic forms, expansion-associated disease has a more generalised atrophy pattern, with thalamic and sometimes cerebellar atrophy as distinguishing features. This should be read probabilistically. It raises coherence between the clinical and genetic hypotheses; it does not by itself establish a molecular diagnosis. The principles and physics of the modalities belong in the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

TRAP **Do not quote one universal prevalence for the repeat expansion in familial FTD.** The New Oxford Textbook of Psychiatry gives about 25% of familial FTD, while separate passages in another major textbook give materially discordant estimates, including an internally lower estimate and another close to the New Oxford figure. State the source and denominator, and acknowledge the discrepancy rather than averaging it or resolving it by fiat.

12.5 When and how to think genetically

Genetic thinking begins with a pedigree, not with indiscriminate testing. Ask about progressive behavioural change, language decline, motor-neurone disease and related neurological syndromes across generations. Clarify the phenotype and age at onset in affected relatives where possible, while recognising that relatives may have been labelled with a psychiatric disorder, “dementia”, weakness, or speech difficulty without a molecular diagnosis. A negative or vague family history lowers confidence but does not convert the case into proven non-genetic disease.

Testing should be considered when **one first-degree relative** is affected in a family compatible with the major progranulin, tau or repeat-expansion causes. Testing for the expansion deserves particular consideration in behavioural-variant frontotemporal dementia, and in late-onset psychiatric presentations when one first-degree relative has frontotemporal dementia or amyotrophic lateral sclerosis. The indication rests on the combined phenotype and pedigree, not on a single symptom.

GOOD TO REMEMBER

Before requesting testing, write down the question the result is meant to answer: confirmation in an affected person, clarification of a familial syndrome, or assessment of an at-risk relative. Then arrange appropriate genetic counselling. A result can affect relatives as well as the person tested, so consent must include familial implications, uncertainty and the possibility that phenotype cannot be predicted precisely.

In practice, construct a multigenerational pedigree with behavioural, language and motor phenotypes recorded separately. When feasible, testing an affected relative first gives the result a clinically anchored context. Use counselling before and after testing, and do not present a pathogenic finding as a precise forecast of symptoms or course. Imaging and pathology associations remain supportive mappings rather than substitutes for a validated molecular result.

Presymptomatic testing raises additional autonomy, confidentiality and family-communication issues. Those issues should be handled through a clinical genetics pathway rather than as a routine laboratory add-on. For the psychiatrist, the core task is to recognise the indication, document capacity and consent carefully, avoid coercive family dynamics, and continue clinical care regardless of whether testing is accepted.

12.6 Motor-neurone disease within frontotemporal degeneration

The overlap is both clinical and pathological. One source reports co-occurring motor-neurone disease in 5-10% of people with FTD; another reports that 10-15% develop it. These adjacent but non-identical ranges should be presented separately because the wording and sampled populations may differ. More importantly, motor disease may precede, accompany, or follow the diagnosis of frontotemporal dementia. Timing therefore cannot be used to exclude the relationship.

Subclinical involvement is even more frequent: electromyographic evidence or mild motor signs such as fasciculations are reported in 30-60% of patients with frontotemporal dementia. Clinically evident motor disease occurs most often with the behavioural variant and less often with semantic or non-fluent aphasic syndromes. Conversely, frontotemporal dementia has been reported in as many as 48% of patients with bulbar-onset amyotrophic lateral sclerosis. This reciprocal enrichment is why screening in only one direction misses disease.

STARTING PRESENTATION	WHAT TO SEEK	INTERPRETIVE POINT
Frontotemporal behavioural syndrome	Weakness, wasting, fasciculations, speech or swallowing change, and evolving functional motor difficulty	Motor disease may emerge before, during or after the cognitive-behavioural diagnosis
Progressive language syndrome	Motor symptoms and signs, particularly when semantic or non-fluent language change coexists with amyotrophic lateral sclerosis	Consider the FTD-MND spectrum, including the kinase loss-of-function association
Motor-neurone disease	Apathy, disinhibition, impaired empathy, executive dysfunction and progressive language change	Cognitive-behavioural symptoms may be part of the same neurodegenerative process
Late-onset psychiatric presentation	Progression, neurological signs, functional decline and a first-degree family history of frontotemporal dementia or amyotrophic lateral sclerosis	A psychiatric label should not terminate neurological and genetic assessment

Pathologically, affected people have ubiquitin-positive and TDP-positive inclusions in lower motor neurons of the spinal cord and in hippocampal granule cells. The shared TDP pathology supports the view that frontotemporal dementia and motor-neurone disease form a spectrum rather than merely co-occurring by chance. This is a disease-model statement, not a claim that every case at either end has clinically obvious involvement of the other.

At review, the motor component should be treated as longitudinal information rather than a single checkbox. Record weakness, wasting, fasciculations, speech and swallowing change, and compare them over time. In the opposite direction, speak with an informant about altered initiative, social conduct, empathy, judgement and language because disability can make cognitive-behavioural change easy to overlook. When examination or history raises concern, neurological collaboration and targeted motor assessment are more defensible than reassurance based on the absence of a dramatic motor presentation.

PITFALL

Do not wait for dramatic weakness before considering the overlap. Mild fasciculations, subtle wasting, changing speech or swallowing, or an unexpectedly abnormal motor examination may be the clue. Equally, do not attribute apathy or reduced speech automatically to respiratory burden or disability in motor-neurone disease without examining executive, behavioural and language change.

12.7 Integrating phenotype, imaging, pathology and pedigree

A disciplined synthesis begins at the bedside and moves outward. First define whether the dominant syndrome is behavioural, semantic-language, non-fluent-language, motor, or combined. Next ask whether the imaging distribution supports a known genetic pattern: symmetric anteromedial temporal involvement for the tau-gene form, asymmetric frontotemporal-parietal

involvement for the progranulin form, or more generalised disease with thalamic and cerebellar involvement for the major expansion. Then add the pedigree and motor examination. Only after these steps should the gene-to-protein map be used to prioritise molecular explanations.

This approach prevents two opposite errors. Overcalling occurs when a suggestive phenotype is treated as proof of a mutation. Undercalling occurs when a psychiatric presentation or apparently isolated motor syndrome is allowed to conceal progressive frontotemporal disease. In both directions, longitudinal change matters: progression and recruitment of another domain make a shared neurodegenerative process more coherent, whereas a static or internally inconsistent picture should prompt reconsideration without reflexively forcing a genetic label.

- **Phenotype:** identify the dominant cognitive-behavioural, language or motor syndrome.
- **Pattern:** describe the atrophy distribution without turning it into a genetic test.
- **Pedigree:** search for differently expressed but potentially related syndromes in relatives.
- **Protein:** map the prioritised gene to tau, TDP, ubiquitin/p62 or FUS pathology.
- **Plan:** arrange counselling and targeted testing when the phenotype-family combination supports it.

Several associations deserve special recall. Progranulin and valosin-containing-protein mutations map to TDP-positive disease. The charged multivesicular-body-protein mutation maps to ubiquitin/p62-positive but TDP-negative disease. The major expansion may show both patterns. The kinase loss-of-function pathway should be considered when amyotrophic lateral sclerosis accompanies semantic or non-fluent aphasia. Together, these exceptions show why “tau-negative” is only an entry point, not a complete molecular diagnosis.

12.8 Writing the high-scoring synthesis

A complete answer can be organised as a compact causal chain. Begin with the pathological division into tau-positive and predominantly TDP-related tau-negative disease, while naming the smaller FUS group. Map the major familial genes to those proteins. Give the characteristic behavioural, language, motor and imaging clues without implying deterministic prediction. Then explain the frontotemporal dementia and motor-neurone disease continuum through shared TDP pathology, reciprocal clinical screening and the major repeat expansion. End with pedigree-based testing and counselling.

Use denominators precisely. “All FTD”, “familial FTD”, “familial FTD-MND”, “amyotrophic lateral sclerosis” and “behavioural-variant FTD” are not interchangeable study populations. The same discipline applies to clinical versus subclinical motor disease. Where authoritative textbook estimates disagree, name the disagreement rather than manufacturing consensus. Where the pathogenic repeat-count threshold is not supplied by the available source, omit the number rather than importing one from memory.

> FTD is a phenotype-to-protein-to-gene problem, and its motor-neurone overlap is best understood as a shared TDP-pathology spectrum.

The safest final sentence is clinical: frontotemporal syndromes require active enquiry for motor disease, motor-neurone disease requires active enquiry for behavioural and language change, and a suggestive family history should trigger genetics-informed assessment rather than an automatic test order. This closes the loop from molecular pathology back to patient care without confusing association with proof.

13 · Differentiating bvFTD from primary psychiatric disorders and from Alzheimer's disease

The difficult bvFTD diagnosis is rarely the florid textbook presentation. It is the patient whose altered conduct, affect, beliefs or social judgement first looks psychiatric, or whose dysexecutive and behavioural syndrome is attributed to Alzheimer disease (AD). The marks lie in showing how the clinician moves beyond resemblance: establish progression, obtain collateral evidence of change from baseline, examine cognition and neurological signs, look for a coherent neurodegenerative pattern, and use disease-directed biomarkers when the competing explanation is AD.

The differential is consequential because psychiatric presentation is common, while neither a behavioural phenotype nor a single abnormal investigation settles aetiology. The task is not to decide whether behaviour is "neurological" or "psychiatric" by appearance. It is to ask which formulation best explains the whole trajectory, including function, cognition, examination, imaging and follow-up. Full diagnostic features of bvFTD and the primary progressive aphasias are covered with the frontotemporal dementia syndromes; this section owns the distinction from primary psychiatric disorders and AD.

13.1 Why bvFTD enters psychiatry first

Behavioural change can precede an obvious complaint of cognitive failure. Patients may therefore be brought for apparent depression, psychosis, bipolar illness, compulsive behaviour or personality change rather than for "dementia." **Behavioural-variant frontotemporal neurocognitive disorder** is often mistaken for a primary mental disorder and commonly first reaches psychiatric services. The psychiatrist is not merely receiving a referral error; psychiatry may be the natural first clinical setting in which the syndrome becomes visible.

The scale of the problem is substantial. **Up to 50%** of people with bvFTD are initially diagnosed as having a primary psychiatric disorder, most commonly depression. A separate source reports that **6%** initially receive a diagnosis of schizophrenia, schizoaffective disorder, bipolar disorder, depression or an unspecified psychotic state. These figures describe differently framed observations and should not be combined into a new estimate. Their shared lesson is that an initial psychiatric label cannot close the case when the subsequent course suggests neurodegeneration.

UP TO 50%	6%	10-40%
FIRST LABELLED PSYCHIATRIC	SPECIFIED INITIAL PSYCHIATRIC LABELS	CLINICAL BVFTD WITH AD PATHOLOGY

Diagnostic revision must nevertheless be disciplined. A psychiatric diagnosis is not disproved merely because the patient later behaves unusually, and a late behavioural presentation is not automatically bvFTD. The correct question is whether progressive neurocognitive decline emerges over time and whether the behavioural presentation sits within convergent evidence of a frontotemporal neurodegenerative syndrome. This approach avoids two forms of stigma: dismissing neurological illness as "just psychiatric," and treating psychiatric illness as neurologically unreal.

RULE **Trajectory outranks first impression.** When bvFTD and a primary psychiatric disorder look similar at presentation, progressive neurocognitive decline and convergent evidence of neurodegeneration carry more diagnostic weight than the original symptom label.

13.2 The longitudinal discriminator

The central discriminator is progressive change from the person's previous level. According to **DSM-5-TR**, the development of progressive neurocognitive difficulties over time helps separate behavioural-variant frontotemporal neurocognitive disorder from primary mental disorders. "Progressive" should not be used as a decorative adjective. It requires evidence that cognition, behaviour or everyday competence is worsening rather than merely remaining unusual, fluctuating with context, or being described differently by different observers.

Collateral history is therefore structural evidence, not an optional supplement. The informant should reconstruct what the person was like before the change, what appeared first, what has accumulated since, and which ordinary roles have become unreliable. Concrete examples are superior to labels. "He has lost empathy" is less informative than a description of repeated socially inappropriate decisions that are new for him and persist despite their consequences. The interviewer then tests whether examples across home, work and relationships form a single evolving pattern.

Serial review also protects against premature certainty. At an early encounter, the clinician may reasonably retain competing formulations. Follow-up asks whether new executive or social-cognitive difficulty appears, whether functional decline becomes demonstrable, and whether neurological or imaging evidence develops. The diagnosis should become more confident because independent lines of evidence converge, not simply because the same behavioural complaint has been repeated at several visits.

The physiological attribution of psychiatric symptoms must also be explicit. If delusions or other symptoms characteristic of a primary mental disorder are judged to arise from the physiological effects of frontotemporal degeneration, the appropriate mental disorder due to frontotemporal degeneration is coded rather than the corresponding primary disorder. The point is causal formulation: the symptom is acknowledged fully, but its diagnostic home changes when the evidence supports a neurodegenerative cause.

TRAP **"Behavioural symptoms mean psychiatric disease; memory symptoms mean dementia."** bvFTD commonly presents to psychiatry, and the useful distinction comes from progression, functional change, focused examination and evidence of neurodegeneration, not from sorting symptoms into crude professional territories.

13.3 Focused assessment when the presentation looks psychiatric

The **Neuropsychiatric International Consortium for Frontotemporal Dementia (NIC-FTD)** developed guidance specifically to assist this differential. Its value is methodological: take a detailed symptom history, add tools designed for frontotemporal presentations, perform a purposeful bedside examination, and integrate the findings with imaging and biological evidence. A generic psychiatric interview alone may describe the presenting symptom accurately while failing to test the neurodegenerative alternative.

FTD-directed instruments supported in that guidance include the Frontal Behavioral Inventory, Stereotypy Rating scale, DAPHNE, and Frontotemporal Dementia versus Primary Psychiatric Disorder Checklist. These instruments organise observations relevant to the differential; they are not substitutes for diagnosis. Their outputs must be interpreted beside informant history, baseline personality and function, because a scale records a pattern within the information supplied to it.

Bedside assessment should deliberately look for **executive dysfunction and impaired social cognition**. Social cognition is typically more impaired in bvFTD than in primary psychiatric disorders. The examiner should also look for parkinsonism and upper motor-neurone signs. These findings matter because they widen the syndrome beyond a purely descriptive account of mood, belief or conduct. Their absence does not settle the differential, but their presence can create clinically coherent convergence.

DIAGNOSTIC AXIS	FINDING THAT FAVOURS BVFTD	INTERPRETIVE CAUTION
Course	Progressive neurocognitive difficulty over time	Do not infer progression from a single dramatic encounter
Function	Continuing decline from the person's baseline	Apparent preservation requires corroboration in real roles
Cognition	Executive and social-cognitive impairment	A general screen may not capture the relevant pattern
Neurological examination	Parkinsonism or upper motor-neurone signs	Absence does not exclude bvFTD
Structural pattern	Frontal or anterior temporal atrophy	Interpret within the clinical syndrome
Neuronal injury marker	Raised serum or CSF neurofilament light may support neurodegeneration	No universal numeric cut-off is supplied here
Follow-up	Accumulating concordant evidence of decline	Non-progression raises the phenocopy question

General administration and scoring of cognitive screens belongs in the Psychometry, Assessment and Descriptive Psychopathology notes. In this differential, the purpose of testing is narrower: identify whether the profile and everyday history support progressive executive and social-cognitive dysfunction, and decide whether specialist neuropsychological assessment would resolve a meaningful uncertainty.

GOOD TO REMEMBER

Write down one concrete baseline-to-current example for behaviour, cognition and function. If the formulation rests only on adjectives such as "odd," "apathetic" or "disinhibited," the evidence is too impressionistic to distinguish degeneration from a primary psychiatric presentation.

13.4 Imaging and biomarkers: seek convergence, not a magic test

Structural evidence becomes useful when it matches the syndrome. **Frontal or anterior temporal atrophy** makes bvFTD more likely in the psychiatric differential. Imaging should therefore be read as part of a pattern: the site of atrophy, the clinical phenotype and the longitudinal course should tell a compatible story. Imaging modality principles and physics are covered in the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

Serum or CSF neurofilament light chain, as an indicator of neurodegeneration, may help distinguish bvFTD from primary psychiatric disorders. It does not identify an FTD subtype, and the permitted evidence gives no numeric decision threshold. A result should therefore be interpreted through the reporting laboratory and the clinical question rather than compared with a remembered value. The absence of a cut-off in this fragment is deliberate, not an invitation to improvise one.

A broader limitation must remain visible: **no serum or CSF biomarker** has been identified that distinguishes FTD from normal controls, and no biomarker differentiates FTD subtypes from one another. By contrast, accepted CSF biomarker combinations can help distinguish AD from FTD. Thus, biomarkers answer different questions. Neurofilament light may contribute evidence of neurodegeneration in the psychiatric differential; the amyloid and tau profile is used when AD is the competing pathology.

The CSF examination procedure and generic constituents are covered in the Neurology foundations for psychiatrists notes. Here, the psychiatrist's task is to choose the diagnostic question, interpret the disease-directed pattern, and avoid turning a supportive result into a stand-alone verdict.

PITFALL

Do not use a normal-looking scan or an isolated laboratory result as a shortcut to "primary psychiatric disorder." Revisit image pattern, timing, functional trajectory and the exact biological question the assay can answer. Conversely, do not diagnose bvFTD from a supportive marker without a compatible syndrome.

13.5 The phenocopy problem

bvFTD phenocopy syndrome describes a subset diagnosed clinically with bvFTD after caregiver-reported personality and behavioural change, but who do not show subsequent disease progression. The key discriminators are the absence of functional decline and the absence of imaging change over time. Its long-term prognosis is good, making an underlying neurodegenerative disorder less likely, although some phenocopy subjects may carry a genetic form of bvFTD.

Phenocopy is not a synonym for fabrication, trivial symptoms or an incompetent initial assessment. It is a longitudinal conclusion arising when a presentation resembles bvFTD but fails to acquire the expected evidence of progression. The clinician must therefore preserve the original observations while revising the causal interpretation. Caregiver reports can be accurate descriptions of changed behaviour even when the subsequent course does not support active neurodegeneration.

The concept prevents diagnostic inertia. Once "bvFTD" appears in a record, later clinicians may reinterpret every behaviour through that label and stop asking whether function or imaging has actually changed. A phenocopy-aware review returns to measurable trajectory. The same discipline works in the opposite direction: absence of decline over an insufficiently informative observation period should not be inflated into certainty. The permitted evidence establishes the defining lack of functional and imaging progression, but does not supply a minimum surveillance interval.

Genetics adds a final caution. The presence of a genetic form in some phenocopy subjects means that non-progression cannot be translated into a universal claim of "non-genetic" or "non-FTD." It is better to document the observed course precisely and keep the aetiological formulation proportionate to the evidence. Detailed FTD molecular genetics belongs with the molecular pathology and motor-neurone-disease overlap discussion.

13.6 Distinguishing bvFTD from AD

Behavioural presentation does not guarantee frontotemporal pathology. Some **10-40%** of patients clinically diagnosed with bvFTD have subsequently shown AD pathology on amyloid PET or post-mortem evaluation. Yet diagnoses made in an expert centre are reported to be almost never due to AD pathology. The contrast is instructive: careful phenotyping and appropriate biomarkers can sharply reduce pathological misclassification, while a label based on behaviour alone remains vulnerable to overlap.

The most relevant clinical competitor is **behavioural-variant Alzheimer disease**. Compared with bvFTD, this AD presentation has worse memory impairment; its memory deficit does not differ from typical AD. Its executive impairment is greater than that seen in both bvFTD and typical AD. Therefore, "executive dysfunction equals bvFTD" is unsafe. The profile must be read across domains, with particular attention to the relative memory burden and the possibility that marked behavioural and executive change can arise from AD pathology.

AXIS	BVFTD	BEHAVIOURAL-VARIANT AD	HOW TO USE THE DISTINCTION
Presenting frame	Behavioural frontotemporal syndrome	Behavioural syndrome caused by AD pathology	Presentation alone does not establish pathology
Memory	Less impaired in the direct comparison	Worse than bvFTD and similar to typical AD	Weight the relative memory burden
Executive function	Less impaired in the reported comparison	More impaired than bvFTD and typical AD	Do not treat executive severity as FTD-specific
Amyloid pathology	Should prompt reconsideration of the clinical label	Supports the AD explanation	Use disease-directed evidence when phenotype overlaps
CSF amyloid signal	Higher relative to AD	Amyloid-beta-42 reduced relative to FTD	Interpret as part of the combined CSF pattern
CSF tau signal	Lower relative to AD	Total tau and phosphorylated tau increased relative to FTD	Combine with amyloid measures and ratios
Final formulation	Integrate phenotype, longitudinal course, imaging and pathology-directed biomarkers		A syndromic label and pathological cause are related but distinct

The accepted **combined CSF amyloid and tau profile** uses amyloid-beta-42, total tau and phosphorylated tau. In AD relative to FTD, CSF amyloid-beta-42 is reduced, while total tau and phosphorylated tau are increased. Diagnostic accuracy can be improved by combined analysis of amyloid-beta-42/amyloid-beta-40, total tau/amyloid-beta-42, or phosphorylated tau/amyloid-beta-42 ratios. No numeric concentration or ratio cut-off is supplied, so the examination-safe answer is the direction and combination, not an invented threshold.

Amyloid PET and CSF biomarkers should be used to answer the pathology question raised by clinical overlap. They do not replace the neurobehavioural history. Conversely, a convincing behavioural history should not prevent pathology-directed testing when the memory and executive profile or other evidence makes AD plausible. Structural and functional neuroimaging in dementia, and fluid biomarkers across dementias, are treated in their dedicated sections; the point here is how they resolve this particular differential.

13.7 Genetic and treatment clues without overreach

Some genetic presentations intensify the psychiatric resemblance. Carriers of **C9ORF72 mutations** may have psychosis as the dominant presenting problem and may initially be diagnosed with obsessive-compulsive disorder, schizophrenia or bipolar disorder. Carriers of GRN mutations may present with bulimia, personality change, sexual disinhibition, ritualistic behaviour or paranoia. These associations should raise diagnostic awareness; they do not justify genetic attribution from phenomenology alone.

Correct differentiation also changes therapeutic expectations. Cholinesterase inhibitors, including galantamine, and memantine are effective for cognitive or behavioural symptoms in AD but have shown no efficacy in FTD. By contrast, **SSRIs and trazodone for bvFTD behavioural symptoms** may improve disinhibition, depression, compulsive behaviour and carbohydrate craving. Olanzapine may help aggression, agitation and psychosis, but dopamine antagonism carries a risk of parkinsonism.

This is a diagnostic consequence, not a full prescribing algorithm. At the LOCUS of care, the clinician still targets the presenting risk and distress while the aetiological formulation is refined. The FOCUS is the symptom causing current impairment, and the MODUS may include environmental, caregiver, psychological and pharmacological measures appropriate to the person. Dementia cognitive-enhancer prescribing and behavioural-symptom management are covered elsewhere in the chapter. General antipsychotic pharmacology and adverse-effect teaching belongs in the Schizophrenia: management, antipsychotics and adverse effects notes.

Response to treatment should not be used backwards as a diagnostic test. Improvement in depression, agitation or compulsive behaviour does not prove a primary psychiatric disorder, because symptom-directed drugs may improve behavioural symptoms of bvFTD. Equally, non-response to an AD-effective cognitive drug does not by itself prove FTD. Treatment history contributes context, while diagnosis continues to rest on syndrome, trajectory and disease-directed evidence.

13.8 An integrated bedside and examination answer

Use a staged formulation rather than a checklist verdict. Begin with the behavioural syndrome and the patient's previous baseline. Establish whether progressive neurocognitive and functional decline is present. Obtain concrete collateral examples. Examine executive function, social cognition and neurological signs. Then seek convergent structural and biological evidence. If there is no functional or imaging progression, reconsider phenocopy. If AD remains plausible, use the memory-executive profile and pathology-directed biomarkers.

MNEMONIC

TRACE the behaviour

Trajectory of progressive decline; **R**elative or caregiver examples from baseline; **A**ssessment of executive function, social cognition and neurological signs; **C**oncordant frontal or anterior temporal imaging; **E**vidence from neurodegeneration markers or AD-directed biomarkers when the differential requires it.

The sequence can be operationalised in a short clinical answer:

- State why bvFTD is frequently labelled psychiatric and identify progressive neurocognitive decline as the key longitudinal discriminator.
- Add focused examination, FTD-directed scales, frontal or anterior temporal atrophy and neurofilament light as convergent evidence, while refusing an unsupported cut-off.

- Close the AD comparison with the memory-executive pattern, amyloid and tau directions, and the possibility of phenocopy when progression fails to appear.

When evidence conflicts, explain the conflict rather than averaging it away. A strongly psychiatric-looking presentation with progressive functional loss, impaired social cognition and concordant atrophy supports bvFTD despite the referral label. A behavioural syndrome with disproportionate memory impairment and an AD biomarker pattern supports behavioural-variant AD. A bvFTD-like presentation without functional or imaging progression raises phenocopy. Each conclusion is a synthesis of independent axes.

> In a behavioural presentation, diagnose bvFTD by progressive neurocognitive decline plus convergent clinical, imaging and biological evidence; then actively exclude phenocopy and Alzheimer pathology.

The mature formulation should state what is established, what remains provisional and what observation would change the diagnosis. That makes uncertainty clinically useful. It also prevents the two classic failures in this topic: retaining a primary psychiatric label after evidence of neurodegeneration has accumulated, and assigning bvFTD from behaviour alone without proving progression or addressing AD overlap.

The treatable, the rapid and the borderline

14 · Reversible dementias: the systematic workup and the treatable causes

Why it matters clinically. The phrase “reversible dementia” is useful only if it changes the diagnostic method. It should prompt a deliberate search for conditions that imitate, cause or compound a neurocognitive syndrome and for which directed treatment may improve cognition, behaviour or function. It should not imply that every positive finding explains the whole presentation, or that treatment restores every patient to the previous baseline.

The marks lie in showing a reproducible sequence: define the syndrome and tempo, identify red flags, perform a focused neurological and general examination, obtain a core laboratory screen, use structural imaging to exclude important intracranial causes, and order targeted tests when the history or examination supplies a reason. The answer becomes clinically mature when it closes the loop from abnormal result to treatment and reassessment rather than ending with a list of investigations.

TEMPO	SCREEN	IMAGE	REASSESS
FRAME THE SYNDROME	CORRECTABLE CAUSES	STRUCTURAL LESIONS	ATTRIBUTE RESPONSE

14.1 What “reversible” should mean

The **reversible-cause paradigm** is an instruction to look for therapeutic opportunity, not a promise of cure. Identifying vitamin deficiency, hypothyroidism or normal-pressure hydrocephalus can produce some symptomatic improvement. That wording matters:

improvement may be partial, a coexisting degenerative disorder may remain, and delay may limit recovery even when the cause is treatable. A candidate who writes “fully reversible” converts a clinical possibility into an unsupported guarantee.

Potentially reversible causes are uncommon but important. In old-age psychiatric and geriatric practice, estimates place them **between 1% and 8%** of patients with dementia. Low frequency does not justify omission because the consequence of missing a treatable intracranial, nutritional, endocrine, infective or metabolic disorder is disproportionate. Conversely, the search must remain disciplined: indiscriminate testing creates incidental findings that can be mistaken for causal explanations.

RULE **Screen systematically, investigate selectively, and verify attribution longitudinally.** A reversible abnormality becomes the cause of cognitive impairment only when the clinical context fits and treatment produces a coherent response, while alternative or coexisting neurocognitive disorders remain under review.

The useful distinction is between reversibility of a cause and reversibility of the entire cognitive syndrome. A deficiency can be corrected without every deficit resolving. A structural lesion can be treated while residual disability persists. A metabolic encephalopathy can clear while an underlying dementia becomes visible. Therefore document the pre-treatment phenotype and function, correct the identified disorder, and reassess the same domains. Without that comparison, “response” becomes impression rather than evidence.

TRAP **Equating a treatable test result with complete reversal of dementia.** The verified outcome is some improvement in symptoms, not a universal return to baseline. The finding may be contributory, coincidental or superimposed on a progressive disorder.

14.2 Build the workup around tempo and examination

A systematic workup begins before the laboratory form. Establish whether the complaint reflects impaired acquisition, retrieval, language, executive control, behaviour, arousal or general slowing, and whether everyday function has changed. Reconstruct onset, progression, fluctuation and the order in which cognitive, gait, urinary, focal neurological and systemic features appeared. An informant can clarify chronology, but discrepancies between accounts should be examined rather than averaged away.

Tempo changes the question. A slowly progressive amnestic syndrome and a fluctuating or rapidly evolving global disturbance should not receive identical differentials. Examine consciousness, attention, speech, eye movements, visual fields, motor tone, power, reflexes, coordination, gait and involuntary movement, then add general examination for nutritional, endocrine, hepatic, renal, infective or malignant clues. The purpose is not to perform an unstructured “complete examination.” It is to generate hypotheses that determine which reversible pathways deserve priority.

Several categories should remain visible from the outset. **Depressive pseudodementia** may resemble a neurocognitive disorder and is differentiated in its dedicated section. **Normal-**

pressure hydrocephalus, subdural haematoma and intracranial tumour are structural possibilities. Deficiency, endocrine, infective, alcohol-related, metabolic and immune-mediated disorders broaden the search beyond the brain scan. The detailed phenotype and management of normal-pressure hydrocephalus and depressive pseudodementia belong to their respective sections of this chapter.

Red flags against a straightforward Alzheimer disease diagnosis include onset at an unusually young age, fluctuation in consciousness, behavioural or personality disturbance overshadowing cognitive loss, rapid progression, gait impairment, lateralised findings, movement disorder, myoclonus, rigidity or bradykinesia. A patient **younger than 65 years** deserves particular diagnostic caution, while development over **6 to 12 months** should redirect attention toward rapidly progressive, infective, immune-mediated and other alternative causes. This section names those routes; rapidly progressive and prion dementias are taught elsewhere in this chapter, and the autoimmune and limbic encephalitis workup belongs in the Neuropsychiatry of systemic and neurological disease notes.

- Describe the cognitive-behavioural syndrome and establish functional decline rather than accepting “memory loss” as the phenotype.
- Use onset, tempo, fluctuation and symptom order to rank structural, systemic and progressive explanations.
- Let examination findings determine targeted tests while retaining a standard core screen for correctable contributors.

14.3 The core laboratory screen

The **baseline laboratory screen** is designed to rule out correctable or contributory causes, not to diagnose a particular degenerative dementia. It includes complete blood count, serum electrolytes, renal and hepatic function, glucose, albumin and protein, vitamin B12 and folate, thyroid-stimulating hormone, urinalysis and a syphilis screen. History-guided additions include a further syphilis test, erythrocyte sedimentation rate and HIV testing. The exact request should reflect local laboratory pathways and the patient’s exposure and examination findings.

WORKUP DOMAIN	CORE OR TARGETED COMPONENTS	QUESTION ANSWERED
Haematological and biochemical	Complete blood count, electrolytes, glucose, albumin and protein	Is a general medical or nutritional disturbance contributing to cognitive dysfunction?
Organ function	Renal and hepatic function	Is impaired clearance or organ failure producing a metabolic cognitive syndrome?
Endocrine	Thyroid-stimulating hormone; calcium is captured through the chemistry assessment	Is an endocrine abnormality present that warrants disease-specific evaluation?
Nutritional	Blood concentrations of vitamin B12 and folate	Is a highly treatable deficiency present?
Infective	Syphilis screening; HIV testing when indicated by history	Does the clinical context justify a directed infective pathway?
Urinary and inflammatory	Urinalysis; erythrocyte sedimentation rate when history guides it	Is there evidence of a contributory systemic process requiring clarification?

Interpretation is the difficult part. An abnormal value must be judged for severity, temporal fit and capacity to explain the observed phenotype. A mild unrelated abnormality should not stop evaluation of a progressive syndrome. Equally, a convincing correctable abnormality should not be filed as “medical” and ignored by psychiatry. The result should trigger confirmation where appropriate, referral or treatment, and repeat cognitive-functional assessment after correction.

Vitamin B12 and folate deficiency are relatively rare but highly treatable causes, so their blood concentrations should be checked in cognitive decline. Numeric reference cut-offs and detailed replacement regimens are intentionally not reproduced here because the disease-level teaching belongs in the Endocrine and metabolic psychiatry notes. The same boundary applies to thyroid, calcium, glucose, hepatic and renal disorders: this chapter teaches their place in the dementia workup, not their full systemic management.

GOOD TO REMEMBER

A laboratory panel is not complete when the results return. Record which abnormality plausibly explains which cognitive or neurological domain, arrange correction or specialist evaluation, and specify what clinical change will be reassessed.

14.4 Structural imaging: routine purpose and selective priority

Structural neuroimaging in the dementia workup looks for normal-pressure hydrocephalus, subdural haematoma, brain tumour or metastases, and cerebrovascular events. Computed tomography is generally sufficient for this exclusionary task. Magnetic resonance imaging is preferred when small or deep infarcts and vascular changes are the concern, and contrast-enhanced magnetic resonance imaging is appropriate when suspicion of tumour is high. The scan answers an anatomical question; it does not replace clinical phenotyping.

Where resources require prioritisation, **selective-imaging criteria** include recent onset, age under the stated threshold, fluctuations, focal symptoms or signs, headache, papilloedema or field defects, malignancy or head injury, seizures, previous stroke or transient ischaemic attacks, incontinence, and gait ataxia or apraxia. Strict application of these criteria detects **about 80%** of potentially reversible causes. That residual miss rate explains why many clinicians prefer structural imaging for every patient with dementia when feasible.

IN INDIA

Computed tomography may be the realistic initial structural study when access, cost or urgency limits magnetic resonance imaging. Use it deliberately to exclude hydrocephalus, subdural collection, mass lesion and major vascular pathology, then escalate to magnetic resonance imaging when the clinical question requires better definition of small, deep or vascular lesions. Resource limitation changes sequencing, not the obligation to recognise red flags.

Imaging should be read against the bedside hypothesis. Ventricular enlargement gains meaning when gait and urinary chronology support hydrocephalus. A collection matters when the history and examination make it plausible. A mass lesion becomes more concerning with focal findings, headache, papilloedema, field loss or malignancy. Conversely, a scan described as “atrophy” does not settle the aetiology of cognitive decline and should not terminate the reversible-cause search.

PITFALL

Do not let a normal scan close the workup or an abnormal scan monopolise it. Structural imaging cannot exclude nutritional, endocrine, infective or metabolic causes, while incidental anatomical abnormalities may not account for the cognitive syndrome.

Imaging modality physics belongs in the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes. Here, the examination answer should remain clinical: state what structural imaging is meant to exclude, why one modality may be preferred for a particular anatomical question, and how imaging findings will be integrated with tempo and examination.

14.5 A cause-based map of treatable pathways

The list of causes is easiest to retain as mechanisms linked to the test that opens each pathway. The categories overlap: a deficiency may coexist with alcohol-related illness, renal or hepatic dysfunction may alter attention as well as memory, and a structural lesion may produce psychiatric or focal neurological features. The map is therefore a set of prompts, not mutually exclusive boxes.

CATEGORY	EXAMPLES TO NAME	HOW THE PATHWAY OPENS	BOUNDARY OF THIS SECTION
Structural and psychiatric	Depressive pseudodementia, benign subfrontal tumour, normal-pressure hydrocephalus, subdural haematoma	Phenotype, neurological examination and structural imaging	Detailed diagnosis and treatment of pseudodementia and hydrocephalus are taught separately
Endocrine and nutritional	Hypothyroidism, hypercalcaemia, hypoglycaemia, Cushing disease and Addison disease; vitamin B12, thiamine, niacin and pyridoxine deficiency	Chemistry, thyroid and vitamin testing guided by clinical context	Disease mechanisms, cut-offs and replacement regimens belong in the Endocrine and metabolic psychiatry notes
Infective	HIV, neurosyphilis and cryptococcal infection; Whipple disease also appears in the broader reversible-cause list	Exposure and systemic clues followed by directed serology or cerebrospinal-fluid evaluation	Use a targeted pathway rather than indiscriminate infective panels
Metabolic and substance-related	Hepatic and renal insufficiency, alcohol-related dementia and **Wilson disease**	Organ-function tests, exposure history and disease-specific evaluation	Systemic diseases themselves belong in the Endocrine and metabolic psychiatry notes
Immune-mediated and endocrine-associated	Hashimoto encephalopathy and paraneoplastic or autoimmune limbic encephalitis	Atypical tempo, behavioural change, seizures or neurological signs prompt specialist assessment	The autoimmune and limbic encephalitis workup belongs in the Neuropsychiatry of systemic and neurological disease notes

Thiamine deficiency requires a specific conceptual pointer. Untreated deficiency may present acutely as Wernicke encephalopathy and may progress to Korsakoff psychosis. The acute entity is the therapeutic emergency; the chronic amnestic sequel is not a reason to delay treatment while awaiting a tidy dementia assessment. Full nutritional-disease teaching belongs in the Endocrine and metabolic psychiatry notes, while Korsakoff amnestic syndrome is covered in the Stroke, TBI, delirium and focal brain syndromes notes.

Alcohol-related cognitive impairment should not be reduced automatically to thiamine deficiency, nor should a vitamin result be treated as an adequate account of every cognitive symptom. The history must link exposure, nutrition, systemic illness and course. This is a general attribution principle across the table: name the treatable route, investigate it appropriately, and continue evaluating any residual pattern after correction.

MNEMONIC

MAP: Metabolic, Anatomical, Psychiatric

Metabolic gathers endocrine, nutritional, infective, organ-failure and substance-related pathways; **Anatomical** prompts hydrocephalus, collection, tumour and vascular disease; **Psychiatric** keeps depressive pseudodementia visible. MAP is a retrieval spine, not a substitute for tempo, examination or targeted testing.

14.6 Targeted infective, metabolic and fluid studies

Neurosyphilis illustrates why "order syphilis serology" is an incomplete answer. Begin with a treponemal test. A negative result makes syphilis very unlikely, although repeat testing may be needed after possible recent exposure. If the treponemal test is positive, add a nontreponemal test. A positive nontreponemal result leads to lumbar puncture and cerebrospinal-fluid analysis; a positive cerebrospinal-fluid VDRL confirms neurosyphilis, while false-negative results can occur. Supportive cerebrospinal-fluid abnormalities may justify presumptive diagnosis and treatment in the appropriate setting, with a low treatment threshold in HIV infection.

The sequence matters because each step changes the probability and the indication for the next. It also prevents two opposite errors: calling any positive syphilis test "neurosyphilis," and excluding neurosyphilis solely because a confirmatory cerebrospinal-fluid test is negative. This is a pointer to the infective reversible pathway, not a full account of general paresis or syphilis treatment.

Hepatic encephalopathy is linked to elevated serum ammonia. The classical examination clue is asterixis, and the electroencephalogram may show triphasic waves. Lactulose acidifies the gastrointestinal tract and converts ammonia to ammonium, which does not cross the blood-brain barrier, thereby reversing the encephalopathy. In a patient labelled as having dementia, fluctuation, motor signs and hepatic dysfunction should therefore trigger a metabolic formulation rather than premature attribution to degeneration.

Cerebrospinal fluid has a selective role in the broader differential. It can demonstrate altered amyloid-beta and tau patterns in Alzheimer disease and elevated **14-3-3 protein** and neuron-specific enolase in prion disease. These studies are not part of a universal "reversible dementia panel," and their interpretation belongs with the suspected syndrome. Dementia-specific biomarkers are covered elsewhere in this chapter; lumbar-puncture technique and generic cerebrospinal-fluid examination belong in the Neurology foundations for psychiatrists notes.

14.7 Acting on a positive screen

A workup earns its value through action. LOCUS depends on clinical stability. A stable patient may complete the core screen and imaging through outpatient care, whereas acute fluctuation, impaired consciousness, rapidly evolving neurological signs or a suspected metabolic or infective encephalopathy requires urgent hospital assessment. The decision is driven by syndrome and physiological risk, not by whether the referral label says "dementia."

FOCUS is phased. The immediate target is a dangerous active cause such as encephalopathy, infection or intracranial pathology. The short-term target is correction of deficiency or endocrine and organ-function disturbance, together with treatment of the identified structural or psychiatric condition. The longer-term target is to determine what cognition and function remain after treatment, whether improvement is stable, and whether a coexisting progressive neurocognitive disorder is becoming clearer.

MODUS includes cause-specific medical treatment, procedural or surgical care for appropriate structural lesions, psychiatric treatment when depressive pseudodementia is established, rehabilitation, risk reduction, family education and serial clinical review. This section can verify lactulose as a route for hepatic encephalopathy and correction of identified vitamin, thyroid or hydrocephalus-related causes as routes to some symptomatic improvement. Detailed regimens, procedures and systemic-disease guidelines belong to their owning chapters or specialist services.

- Confirm that the abnormality is real and clinically relevant before declaring it causal.
- Treat or refer according to urgency, while documenting the cognitive, neurological and functional baseline.
- Reassess the same domains after correction and continue the dementia evaluation if deficits persist or progress.

Improvement should be described by domain rather than as a global impression. A patient may become more alert while memory remains impaired, walk better while executive dysfunction persists, or regain initiative while instrumental function remains limited. Domain-based reassessment prevents both therapeutic nihilism and exaggerated claims of cure. It also provides a cleaner baseline for later follow-up.

14.8 Examination synthesis

A strong long answer opens by qualifying the term: potentially reversible causes are uncommon, clinically important, and may yield some improvement rather than guaranteed full recovery. It then presents the sequence of history, tempo, cognitive and neurological examination, core laboratory tests and structural imaging. Red flags justify accelerated or targeted investigation, while the positive result must be linked to a plausible mechanism and followed through to treatment and reassessment.

The cause list should be organised, not dumped. Name structural and psychiatric mimics, endocrine and nutritional causes, infections, organ-failure and substance-related states, and immune-mediated disorders. Add one or two pathways to show method: the staged serological and cerebrospinal-fluid approach to neurosyphilis, or the clinical and treatment logic of hepatic encephalopathy. Then state the ownership boundaries rather than drifting into full endocrine, nutritional, autoimmune or hydrocephalus teaching.

The final sentence should restore diagnostic humility. Correction of a reversible contributor neither proves that it caused every deficit nor excludes a coexisting progressive dementia. Failure to improve does not retrospectively make the workup unnecessary, because the investigation may

still have treated a harmful contributor. The defensible endpoint is a revised formulation based on treatment response, residual phenotype and longitudinal course.

> "Reversible dementia" means a systematic search for treatable contributors followed by cause-specific treatment and domain-based reassessment; it does not mean that complete cognitive recovery is guaranteed.

15 · Depressive pseudodementia: features and differentiation from true dementia

Why it matters clinically. A patient with low mood, slowed thinking and poor recall may look as though a progressive dementia has begun. The reverse error is equally consequential: a progressive cognitive disorder may first become visible through depression. The marks in this topic therefore lie not in reciting a stereotype, but in showing how history, observed behaviour, repeated bedside testing, mood assessment and longitudinal review are integrated without prematurely declaring either reversibility or degeneration.

Depressive pseudodementia is the more common and clinically important form of pseudodementia in which depressive illness presents with prominent memory, concentration and intellectual complaints. The traditional label remains useful as a clinical shorthand, but it can falsely imply that the impairment is unreal, voluntarily produced or guaranteed to disappear. When the cognitive disturbance meets the clinical threshold for dementia, many neuropsychiatrists prefer **reversible dementia syndrome**. This wording recognises genuine disability while leaving the trajectory open. The central task is to identify the depressive contribution, treat it, and continue surveillance for an underlying progressive disorder.

15.1 Terminology, validity and the continuum model

"Pseudo" should never be read as "pretended." The patient may be unable to work, manage domestic tasks or sustain ordinary conversation because attention, motivation, speed and retrieval have collapsed under a severe depressive state. The cognitive syndrome is real even when it later improves. A more accurate conceptual model is a **clinical continuum**. At one end, the depressive syndrome accounts for most of the cognitive disturbance. At the other, a progressive subclinical neurodegenerative disorder generates the impairment and depression adds to its severity. Many patients cannot be placed confidently at either end during the initial assessment.

This continuum changes the diagnostic question. The clinician should not ask only, "Is this depression or dementia?" The more useful questions are: How much of the current disability is state-dependent? Is there evidence of a progressive cognitive process independent of mood? What happens when time, encouragement and repeated learning opportunities are provided? Does cognition improve with remission, remain impaired, or deteriorate despite mood recovery? These questions allow a provisional formulation that can be revised as longitudinal evidence accumulates.

RULE **Reversibility is a longitudinal observation, not an assumption made from the first interview.** Diagnose and treat the depressive syndrome, document the cognitive phenotype, investigate competing causes, and deliberately reassess cognition after mood improvement.

The older description of pseudodementia as a readily correctable cause is therefore incomplete. It remains treatable, but it may also be a **harbinger of dementia**. The label should open a treatment window without closing the diagnostic work-up.

15.2 Clinical phenotype of depressive pseudodementia

The usual presentation is an older adult who complains of poor memory and concentration, sometimes while denying overt depression. The onset is relatively acute or subacute rather than an unmistakably gradual story of accumulating decline. Patients may emphasise inability, mental blankness and loss of competence. Their account can sound globally catastrophic, yet performance across the encounter may fluctuate. A past or family history of affective illness supports the formulation, but it cannot prove that the cognitive syndrome is purely depressive.

When sadness is minimised or denied, look for the **biological and cognitive features of depression**: disturbed sleep, low energy, psychomotor retardation, loss of interest in usual activities or hobbies, pessimistic thinking and rumination. The distinction is not made by vegetative symptoms alone. Sad mood and psychic depressive features are particularly useful when deciding whether depression is superimposed on a dementia syndrome rather than dementia occurring alone. The interview should therefore examine affective experience, self-appraisal and loss of pleasure as carefully as appetite or sleep.

MNEMONIC

PACE

A bedside spine for depressive pseudodementia: **P**resentation relatively acute or subacute; **A**ffective and biological clues; **C**ognitive performance patchy; **E**ncouragement and repeated trials may improve performance.

Psychomotor slowing deserves active observation. Long response latency, sparse speech and reduced spontaneous elaboration may look like aphasia or loss of knowledge. In depressive pseudodementia, however, language output is often slow and limited without the paraphasic errors that would make an acquired language disorder more persuasive. The examiner should separate reduced production from disordered language content. The same principle applies to everyday function: failure to initiate an activity is not equivalent to loss of the cognitive capacity needed to perform it, although both can coexist.

15.3 Bedside cognitive pattern

Attention is commonly impaired, so later failures may reflect weak engagement with the material at the point of encoding. Registration of a name and address and performance on digit span may initially be poor. The crucial manoeuvre is not merely to record failure, but to repeat the task with

calm pacing, sufficient time and explicit encouragement. Improvement across repeated trials supports a depressive contribution. A single hurried score conceals this dynamic information.

Memory and executive performance are often **patchy and inconsistent**. A patient may fail a formal item, then later reveal relevant information in conversation; solve a demanding component, then stop at an easier component; or show wide variation when effort and confidence change. Inconsistency is useful only when observed carefully. It should not be translated into malingering, and it does not exclude coexisting brain disease.

The characteristic verbal response is the **"don't know" answer**. This differs clinically from an incorrect but confident response, circumlocution or a paraphasic error. The depressed patient may abandon retrieval early, express certainty that recall is impossible, or refuse to guess because failure fits a pessimistic self-model. Encouragement may recover information that was accessible but not produced. By contrast, persistent loss of newly learned material despite adequate acquisition raises greater concern for a progressive amnestic process. The relevant contrast is with rapid forgetting, which is not the expected depressive pattern.

Acute	Patchy	"Don't know"	Recheck
ONSET CLUE	PERFORMANCE	RESPONSE STYLE	TRAJECTORY

Testing conditions themselves become diagnostic evidence. Depression need not substantially impair performance when the examiner provides adequate time and encouragement. This **effort and encouragement effect** is not a licence to coach the patient toward a desired diagnosis. Standard instructions should be preserved, while latency, persistence, response to reassurance and learning over repeated attempts are documented separately. The examiner is studying the process of performance, not only the final total.

15.4 Differentiation from a progressive dementia syndrome

The accompanying figure and the table below organise the discriminating features. No feature is decisive in isolation. Their value comes from convergence: tempo, depressive phenomenology, response style, consistency, learning across trials, error type and subsequent course should point in the same direction before confidence rises.

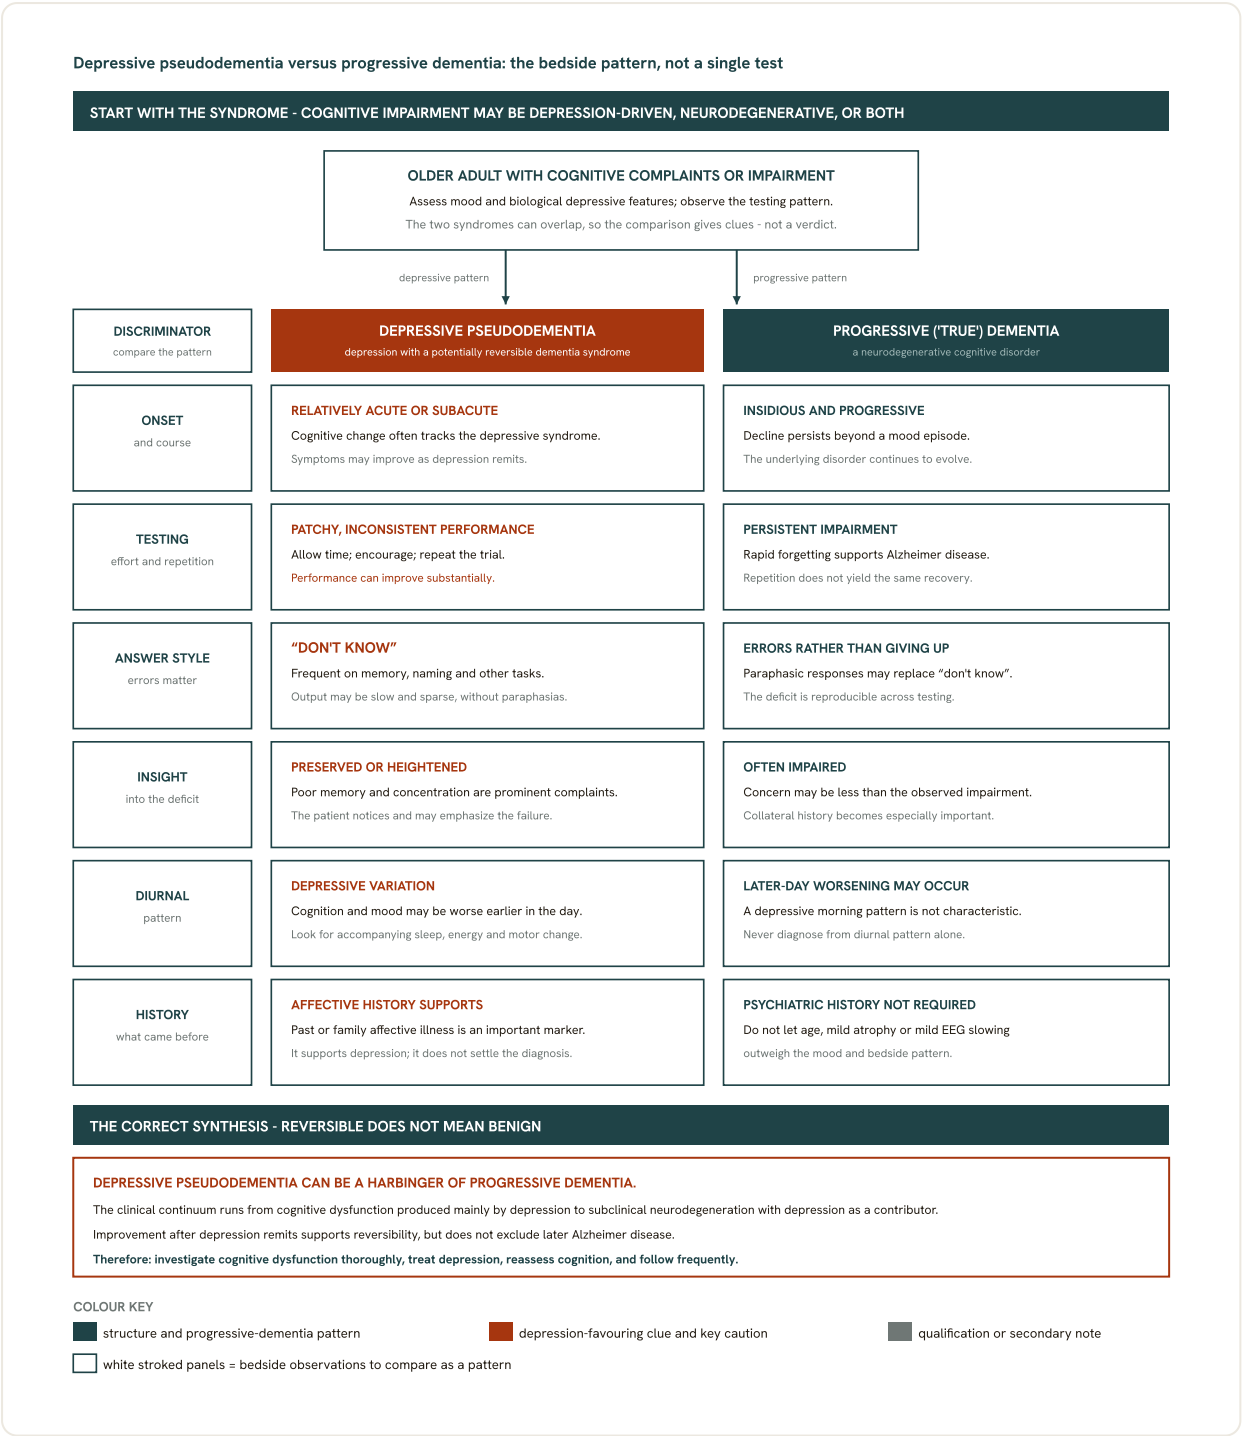


Fig 32.6 **Depressive pseudodementia versus progressive dementia.** Relatively acute onset, patchy effort-dependent performance, frequent "don't know" responses, preserved concern about deficits, depressive diurnal variation and an affective history favour depression with a reversible dementia syndrome; insidious progression, rapid forgetting, reproducible errors and reduced insight favour progressive dementia. The syndromes overlap: apparent cognitive recovery after depression does not exclude later Alzheimer disease, so diagnostic workup and frequent follow-up remain essential. (Kaplan & Sadock's Comprehensive Textbook of Psychiatry, 2024; Kaufman's Clinical Neurology for Psychiatrists; Hodges, Cognitive Assessment for Clinicians.)

CLINICAL AXIS	DEPRESSIVE PSEUDODEMENTIA FAVOURED	PROGRESSIVE DEMENTIA FAVOURED	HOW TO USE THE DISTINCTION
Tempo	Relatively acute or subacute onset	Progressive course that becomes clearer with serial history	Anchor the history to change from the patient's prior level rather than to the consultation date
Subjective account	Prominent complaints of poor memory and concentration; overt depression may be denied	Complaint and observed impairment may not align	Obtain an informant account without treating either account as automatically superior
Depressive clues	Sleep disturbance, low energy, psychomotor slowing, anhedonia, pessimism and rumination	Depressive features may be absent, secondary or superimposed	Assess psychic depressive experience as well as vegetative change
Task engagement	Early abandonment, low confidence and "don't know" responses	May attempt answers despite errors	Record the response process rather than reducing it to correct versus incorrect
Consistency	Patchy performance across tasks or across the interview	A more coherent cognitive pattern may emerge	Seek convergence across history, examination and formal assessment
Repeated trials	Registration and span may improve with repetition, time and encouragement	New learning may remain impaired despite adequate opportunity	Document acquisition and retention separately
Error quality	Slow, sparse language without paraphasic errors; naming may elicit "don't know"	Specific error patterns may suggest an acquired cognitive syndrome	Distinguish reduced output from disordered content
Course after treatment	Cognition may improve or largely subside with remission, although deficits may remain	Persistent or progressive impairment becomes more evident	Reassess formally after mood improvement and continue follow-up

The table is asymmetric by design: the permitted bedside evidence describes the depressive pattern in detail, but "true dementia" is not a single neuropsychological phenotype. Its presentation varies with aetiology and disease stage. Accordingly, the right-hand column states directional concerns rather than pretending that every progressive dementia produces the same errors. Detailed dementia phenotyping belongs with the disorder-specific sections and formal cognitive assessment.

15.5 Diagnostic method and common clinical errors

A safe assessment begins with chronology and functional change. Establish what changed, how quickly it became apparent, whether the course tracks mood, and which activities fail because the patient cannot initiate them versus cannot understand, remember or sequence them. Interview an informant where possible, but examine discrepancies rather than resolving them by authority. Observe affect, psychomotor activity, spontaneity, latency and error style throughout the encounter. Bedside tasks then test hypotheses raised by the history.

The working formulation should remain layered. State whether a depressive syndrome is present, describe the cognitive domains affected, record features supporting state-dependent impairment, and state whether a progressive disorder remains possible. In an older depressed adult, cognitive dysfunction warrants a thorough diagnostic work-up and frequent follow-up aimed at identifying treatable neurocognitive disorders. The systematic reversible-dementia work-up is covered elsewhere in this chapter. General administration and scoring of cognitive screens belong to the Psychometry, Assessment and Descriptive Psychopathology notes.

- Repeat selected learning and attention tasks rather than treating the initial failure as final.
- Record “don’t know,” paraphasic errors, intrusions, abandonment and response latency separately.
- Link functional loss to its mechanism: poor initiation, poor persistence, failed acquisition or failed retention.
- Arrange reassessment after meaningful mood improvement, even when cognition appears to recover.
- Escalate to psychiatric opinion and formal neuropsychological evaluation when simple cognitive testing cannot resolve the distinction.

> **PITFALL**
>
> Age-related mild slowing on electroencephalography or cerebral atrophy on structural imaging can be overvalued when mood is under-assessed. Such findings do not substitute for a clinical account of tempo, depressive phenomenology, testing behaviour and trajectory. Imaging interpretation for dementia belongs with the dementia-specific assessment; modality principles are reviewed in the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

At times, simple cognitive tests cannot distinguish depressive pseudodementia from a progressive dementia. That is not a failed examination. It is the correct recognition that bedside certainty has reached its limit. Psychiatric assessment clarifies the depressive syndrome, while formal neuropsychological evaluation examines the architecture and validity of performance in greater depth. Longitudinal review then supplies evidence that no cross-sectional test can provide.

15.6 The remission paradox

The classic teaching is that cognition improves or largely subsides after remission of depression, hence the terms **dementia syndrome of depression** and depression with reversible dementia.

This is clinically useful because it directs active treatment and reassessment. Yet "reversible" must describe an observed course, not a diagnostic promise. Even after mood symptoms settle, abnormal executive function, processing speed, memory consolidation and working memory may persist in many older patients with depression. Older adults with both depression and diabetes are particularly vulnerable to cognitive deficits.

TRAP **Do not write that late-life depression-associated cognitive impairment either always reverses or generally never reverses.** Separate chapters in the same major textbook place different emphasis on the evidence: the geriatric account defines a syndrome that improves or largely subsides with remission, while the main dementia discussion stresses persistence, especially in memory and executive function. The defensible answer surfaces the tension: substantial improvement may occur, residual deficits are common, and subsequent decline remains possible.

This apparent contradiction dissolves clinically when state and trait are separated. Depressive slowing, reduced effort and inefficient encoding may improve, while pre-existing vulnerability or an emerging progressive process remains. The patient can therefore be better without being cognitively normal. A binary endpoint such as "pseudodementia cured" misses domain-specific residual deficits, and a stable total score may conceal a changed pattern of strengths and weaknesses.

Remission also provides a cleaner diagnostic window. Reassessment should ask whether attention and engagement improved, whether acquisition now benefits from repetition, whether functional independence returned, and whether any impairment continues to progress outside depressive episodes. The goal is not to retrospectively prove that the original diagnosis was wrong. It is to update the formulation using better conditions for observing cognition.

15.7 Prognosis and surveillance

The prognostic message is deliberately uncomfortable. Pseudodementia predicts later **Alzheimer disease** in some, but not all, patients. Depression may have been a prodromal manifestation of the primary cognitive disorder, or the depressive cognitive syndrome may identify a subgroup with an underlying progressive process. Apparent recovery therefore lowers immediate impairment but does not erase future risk.

Even among people with nearly complete cognitive recovery, irreversible dementia has been reported at about **20% per year** during follow-up. Many may remain nondemented for **1 to 2 years** before irreversible dementia appears. In older adults with depression and mild cognitive impairment, about **40%** progressed to that disease within **27 months**. Recency adds prognostic weight: active depression in the recent past is a marker of greater risk than a remote history. This is another reason to interpret the cognitive presentation in temporal relation to the depressive episode rather than treating "past depression" as a single undifferentiated variable. Another account reports that patients initially appearing to have depression-induced cognitive impairment develop the same dementia **almost six times** as frequently as age-adjusted counterparts.

GOOD TO REMEMBER

A good response to depression treatment changes care but does not end it. Record a post-remission cognitive baseline, explain the need for surveillance without deterministic language, and ask the patient and family to return if function declines. Follow-up is part of the diagnosis because conversion risk remains elevated even after apparent cognitive recovery.

These figures should not be converted into a prediction for an individual patient. Their proper use is to justify structured follow-up, not to announce inevitable degeneration. The clinician should avoid both false reassurance and fatalism. A patient may remain cognitively stable; another may reveal a progressive disorder only after depressive symptoms remit. Clear documentation of baseline function, mood state and cognitive pattern makes later change interpretable.

15.8 The dissociative look-alike

The umbrella term pseudodementia has also been used for **hysterical pseudodementia**, a distinct syndrome from depressive pseudodementia. Its onset may be fairly abrupt, and the patient often appears unconcerned. Memory loss can be paradoxically greater for highly salient personal or remote autobiographical events than expected in an organic amnesic disorder; personal identity may be lost. Performance may become strikingly worse under formal testing than during informal conversation about recent events.

An identifiable precipitant such as bereavement, marital difficulty or offending, together with a past psychiatric history, may support this formulation. Some patients show **Ganser syndrome** features, particularly approximate answers: the response comes close to the correct answer but remains incorrect. This pattern differs from the pessimistic “don’t know,” slowed sparse output and repetition-sensitive acquisition that characterise depressive pseudodementia.

The distinction matters because “pseudodementia” is not a unitary explanation. Abrupt onset and inconsistent testing can occur in both forms, but the affective context, autobiographical pattern, concern about deficits and quality of wrong answers differ. Neither pattern should be diagnosed from theatricality or intuition. The clinician must describe positive features, assess mood and context, and still exclude an underlying neurocognitive disorder.

15.9 Examination synthesis

A high-quality answer begins by defining depressive pseudodementia as genuine cognitive impairment associated with depression, preferably framed as a potentially reversible dementia syndrome rather than feigned illness. It then states the continuum from predominantly depressive cognitive dysfunction to depression superimposed on subclinical progressive disease. The discriminating core follows: relatively acute or subacute onset, depressive biological and psychic features, prominent complaints, impaired attention, patchy performance, “don’t know” responses, slow sparse language without paraphasia, and improvement with time, encouragement and repeated trials.

The answer must then qualify the distinction. Bedside testing can remain indeterminate, requiring psychiatric and formal neuropsychological assessment. Improvement with remission may be large, yet executive, processing-speed, consolidation and working-memory deficits can persist. Finally, apparent recovery does not abolish later dementia risk, so thorough work-up and repeated follow-up are integral rather than optional. This final qualification separates an old mnemonic answer from clinically current reasoning.

> Depressive pseudodementia may improve with remission, but its bedside pattern is probabilistic and its future course is not guaranteed: treat the depression, document the cognitive process, and follow the patient longitudinally.

16 · Normal pressure hydrocephalus: triad, diagnosis and shunt response

Why it matters clinically. Normal-pressure hydrocephalus is a high-value diagnosis because an apparently dementing illness may reflect a disorder of cerebrospinal-fluid dynamics for which drainage can produce meaningful clinical relief. The marks lie less in reciting a slogan than in recognising the internal order of the syndrome: gait disturbance leads the clinical picture, cognition has a subcortical character, urinary symptoms evolve, imaging supports but does not independently establish the diagnosis, and response to temporary cerebrospinal-fluid removal helps select patients for permanent drainage.

The central reasoning problem is uncertainty. Ventricular enlargement is not unique to this disorder, the complete syndrome need not present as a perfectly synchronised package, a favourable drainage test is more persuasive than an unfavourable one, and shunt benefit is not completely predictable. A strong answer therefore moves from phenotype to concordant investigations and then to a balanced therapeutic decision. It neither dismisses a treatable possibility because cognition is impaired nor labels every enlarged ventricular system as the cause of that impairment.

GAIT	COGNITION	BLADDER	DRAINAGE
LEADING CLUE	SUBCORTICAL	PROGRESSIVE	PREDICTIVE TEST

16.1 The syndrome and the pressure paradox

Normal-pressure hydrocephalus describes ventricular enlargement occurring despite apparently normal cerebrospinal-fluid pressure. The apparent paradox matters. “Normal pressure” should not be interpreted as proof that pressure has never been abnormal or that cerebrospinal-fluid dynamics are physiologically normal. Intermittent pressure surges may contribute even when pressure appears normal at examination. The diagnostic target is therefore a dynamic clinical syndrome, not merely a single pressure reading.

Hydrocephalus is classified according to the site of obstruction to cerebrospinal-fluid flow. In **communicating hydrocephalus**, the block lies outside the ventricular system; in **non-communicating hydrocephalus**, it lies within the ventricular system. Compensated hydrocephalus differs conceptually because clinical signs and cerebrospinal-fluid dynamics have

stabilised while pressure remains elevated. These distinctions prevent the loose use of "hydrocephalus" as though every form shared the same pressure state, mechanism, presentation and management logic.

RULE **Diagnose a clinicoradiological syndrome, not a scan or a pressure value in isolation.** The diagnostic argument becomes coherent only when the characteristic clinical pattern, ventricular enlargement and the response to temporary drainage are interpreted together.

For the psychiatrist, the syndrome sits at the interface between cognitive medicine and neurological examination. A referral framed as "dementia" may conceal the decisive motor clue, while a referral framed as "walking difficulty" may understate executive slowing or loss of initiative. History and examination must therefore be deliberately integrated. The family's account of sequence is especially valuable because the order in which gait and cognitive symptoms appeared changes the weight placed on this diagnosis.

16.2 The triad as an ordered clinical pattern

The **clinical triad** comprises gait apraxia, cognitive impairment and urinary incontinence. Its components are not equally informative at every stage. Gait apraxia is generally the initial, most consistent and most prominent feature. The cognitive disorder follows a subcortical pattern, while urinary urgency and frequency may precede established incontinence. Thus, the triad is best used as a structured search pattern rather than as a demand that every component be fully developed before the diagnosis is considered.

DOMAIN	CHARACTERISTIC EXPRESSION	DIAGNOSTIC USE	EXPECTED TREATMENT SIGNAL
Gait	Apraxic, magnetic pattern with impaired alternation, weight transfer and turning	Usually the leading and most dependable clinical clue	Often the earliest domain in which improvement becomes evident
Cognition	Slowed thought, reduced initiative and relative preservation of language skills	Supports a subcortical cognitive syndrome rather than an isolated cortical-language presentation	May not improve during a temporary drainage test even when gait improves
Urinary function	Urgency and frequency progressing to incontinence	Evolution matters more than waiting for complete loss of continence	Can improve after treatment

The ordering has a practical consequence. If cognitive decline clearly precedes gait dysfunction, normal-pressure hydrocephalus becomes less likely. This is not equivalent to saying that cognitive symptoms exclude it. Rather, the chronology is discordant with the usual pattern and should prompt a more critical search for another explanation of cognitive decline, ventricular enlargement, or both. Conversely, a prominent gait disorder accompanied by compatible subcortical cognitive and urinary changes creates a more persuasive syndrome.

MNEMONIC

Gait leads, thought slows, bladder warns

This spine preserves the clinically useful order without importing an unverified eponym or ward mnemonic: gait is the leading clue, cognition is slowed and subcortical, and urinary urgency or frequency warns before established incontinence.

16.3 Gait apraxia: the bedside discriminator

Gait apraxia is the key bedside feature. The problem is not captured adequately by saying only that the patient "walks slowly." The patient may fail to alternate leg movements correctly, fail to transfer weight onto the forward foot, or pick up the same leg again instead of progressing in sequence. The weight-bearing foot may appear stuck to the floor, producing the magnetic quality. Turning is fragmented and requires multiple small steps rather than a smooth change of direction.

The examination should make these failures visible. Observe initiation, straight walking, foot clearance, alternation, weight shift and turning as separate tasks. A global label such as "unsteady" loses the information needed for diagnostic reasoning. Ask what exactly interrupts forward progression: failure to release the planted foot, failure to move weight forward, disordered sequencing, or instability during the turn. Description precedes interpretation. This is particularly important when the history has already been framed as falls, weakness, ageing or "poor effort."

A useful bedside contrast is that stepping over a stick may remain preserved. This indicates that the stepping reflex is relatively intact even though ordinary gait sequencing is impaired. The contrast helps clarify why the disorder is described as apraxic: the legs can produce a stepping response under a particular cue, but the self-generated sequence of normal walking is disrupted. It also guards against treating the abnormality as a simple loss of strength.

- Watch the first steps after the command to walk, because initiation may expose the magnetic quality.
- Describe forward weight transfer and foot alternation, and observe a turn rather than judging only straight walking.
- Use the preserved ability to step over an obstacle as a clinical contrast, not as a reason to call the gait normal.

GOOD TO REMEMBER

Record gait before and after temporary cerebrospinal-fluid removal in concrete, observable language. "Better" is weaker evidence than a documented change in initiation, alternation, weight transfer, foot release or turning.

Gait is also the most useful early treatment signal. It is generally the first feature to improve, and improvement after temporary cerebrospinal-fluid removal supports the diagnosis and predicts benefit from permanent drainage. Cognitive change need not accompany that immediate gait

response. An examiner therefore expects the candidate to prioritise gait assessment both for recognition and for testing reversibility.

16.4 Cognitive and urinary components

The dementia has a **subcortical cognitive pattern**. Thought is slowed, and the motor slowing may be evident alongside it, while language skills are relatively spared. This profile matters because "dementia" is otherwise too broad a label to connect the cognitive presentation to the gait syndrome. The psychiatrist should seek reduced speed of thinking, delayed generation of responses and diminished spontaneous engagement rather than assuming that every cognitive complaint will present primarily as forgetting or loss of language.

Abulia is common. Clinically, diminished initiative may be misread as indifference, depression, oppositional behaviour or poor cooperation. Within the syndrome, it should be interpreted beside slowed thought and gait rather than in isolation. The interview may appear sparse because the patient generates little spontaneously, yet this does not by itself establish a primary mood explanation. Depressive pseudodementia and broader comparative differentiation from depression and delirium are addressed elsewhere in this chapter; the present task is to recognise the characteristic cognitive tone of normal-pressure hydrocephalus.

The urinary component begins with **urgency and frequency** and progresses to incontinence. Asking only whether the patient is incontinent can therefore miss the evolving syndrome. History should establish the direction of change and its relationship to gait and cognition. Urinary symptoms also have therapeutic relevance because they can improve with treatment. Their presence strengthens the clinical pattern, but the leading diagnostic weight remains with gait.

TRAP **Waiting for fully developed incontinence before considering the diagnosis.** Urgency and frequency may be the urinary expression before incontinence appears, while gait usually supplies the earlier and more reliable clue.

16.5 Imaging, cerebrospinal fluid and diagnostic synthesis

Structural imaging shows **ventricular dilatation**, particularly involving the temporal horns. There may also be signs suggesting cerebrospinal-fluid reabsorption across ventricular surfaces. These findings are supportive, but they are nonspecific. Imaging alone cannot reliably establish that ventricular enlargement is responsible for the patient's gait and cognitive syndrome.

The central imaging differential is **hydrocephalus ex vacuo**, in which ventricular enlargement resembles that seen with cerebral atrophy. The practical question is causal: does the enlarged ventricular system belong to a pressure-dynamics syndrome capable of responding to drainage, or does it accompany loss of brain volume? A report of "dilated ventricles" does not answer that question. Clinical sequence and phenotype remain indispensable, and temporary drainage may add functional evidence.

Cerebrospinal-fluid pressure is normal in normal-pressure hydrocephalus, as are protein and glucose concentrations. Normal routine cerebrospinal-fluid findings therefore do not invalidate

the syndrome. Equally, the word “normal” must not become a shortcut to diagnosis: normal pressure and constituents do not distinguish this condition without the compatible clinical and imaging context. The procedure and generic study of cerebrospinal-fluid constituents belong in the Neurology foundations for psychiatrists notes.

Imaging modality principles and physics belong in the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes. Here, the clinically relevant use of imaging is narrower: establish ventricular enlargement, note supportive features, examine whether atrophy offers a competing explanation, and integrate the result with gait-led clinical findings. This preserves the distinction between knowing how an imaging modality works and knowing what a finding contributes to a particular diagnosis.

PITFALL

Do not let the radiology label close the case. Ventricular enlargement may resemble hydrocephalus ex vacuo, so attributing cognition and gait to normal-pressure hydrocephalus from imaging alone risks both inappropriate surgery and failure to identify another cause.

A disciplined synthesis can be stated as a chain. First, establish whether the gait phenotype is genuinely apraxic and whether it led the illness. Next, seek the compatible subcortical cognitive and evolving urinary features. Then ask whether imaging demonstrates ventricular enlargement without treating that finding as diagnostic by itself. Finally, use the response to temporary cerebrospinal-fluid removal to refine the probability of shunt-responsive disease. Each link answers a different question; none is a substitute for the others.

16.6 Temporary cerebrospinal-fluid removal

The **large-volume tap test** involves withdrawing **30-60 mL** of cerebrospinal fluid by lumbar puncture. An alternative is a series of **three** lumbar punctures. The test is worthwhile because it asks a functional question that static imaging cannot answer: does temporary drainage improve the characteristic clinical deficit?

Gait is the principal outcome. Improvement in gait after cerebrospinal-fluid removal indicates normal-pressure hydrocephalus and predicts benefit from permanent drainage. Dementia need not improve during the test for the result to be meaningful. This asymmetry is clinically sensible within the verified syndrome: gait is usually the leading feature and the first to respond to treatment. Insisting on an immediate cognitive response would therefore discard useful positive evidence.

- Define gait abnormalities before drainage, then re-examine the same domains after drainage.
- Interpret gait improvement as positive evidence for both the diagnosis and likely benefit from permanent drainage.
- Do not require cognitive improvement for a positive gait response, and do not use a negative test as an absolute exclusion.

A negative test does not preclude the diagnosis. That statement is more than a caveat: it defines the direction in which the evidence should be used. A positive gait response is affirming and predictive; a negative response is not sufficiently decisive to close the diagnosis. When the phenotype and imaging remain persuasive, the result must be placed back into the total clinical argument rather than converted into a binary rule.

General lumbar-puncture technique, contraindications and the generic cerebrospinal-fluid examination procedure belong in the Neurology foundations for psychiatrists notes. Within this topic, the examiner is testing the indication, volume of removal, outcome domain and interpretation. The best answer therefore centres on large-volume removal, documented gait response, its predictive value for shunting, and the limited exclusionary value of a negative result.

16.7 Shunt treatment and response

Management is a selection problem. LOCUS begins with coordinated specialist assessment and proceeds to neurosurgical care when permanent drainage is being considered. FOCUS is functional: relieve the gait-led syndrome, improve urinary dysfunction where possible, assess cognitive change without demanding it as the earliest response, and avoid exposing an unsuitable patient to a procedure whose benefit cannot be predicted with certainty. MODUS is procedural drainage supported by careful clinical observation and follow-up; the verified treatment is insertion of a ventricular shunt.

A **ventricular shunt** is inserted into a lateral ventricle and drains cerebrospinal fluid into the chest or abdominal cavity. This can relieve normal-pressure hydrocephalus. The decision is not trivial, because a clinically beneficial response is difficult to predict. A compatible gait-led syndrome and improvement after temporary drainage provide a more persuasive basis for treatment than ventricular enlargement alone.

Response should be understood by domain. Gait is generally the first feature to improve. Urinary symptoms can also improve. Cognitive improvement may be less immediately evident during temporary drainage and should not be made the sole test of success. This domain-based view prevents the vague statement that “the patient improved” and keeps assessment aligned with the syndrome’s natural hierarchy.

Shunt insertion is prone to complications, including subdural haematoma and infection. Surgical complications can be serious and occur in up to **30%** of cases. The procedure can be beneficial in hydrocephalus, but it should not be undertaken lightly. The therapeutic argument must therefore contain both halves: a potentially reversible disorder deserves active assessment, and the possibility of reversibility does not justify indiscriminate shunting.

The treatment discussion should link the evidence explicitly. A gait-led clinical syndrome supplies pre-test coherence. Imaging confirms ventricular enlargement but remains nonspecific. A positive gait response to cerebrospinal-fluid removal supports normal-pressure hydrocephalus and predicts benefit from permanent drainage. A negative response weakens neither the diagnosis nor the treatment case absolutely. The final decision requires judgement because neither imaging nor temporary drainage makes shunt response perfectly predictable.

No single numerical improvement rate should be quoted when it is not supported by the source base. The defensible examination statement is that benefit is difficult to predict, gait is the earliest expected responder, urinary symptoms may improve, and serious complications must be weighed. This is both more accurate and more clinically useful than false precision.

16.8 A compact diagnostic answer

A complete short-note answer begins by defining normal-pressure hydrocephalus as ventricular enlargement despite apparently normal cerebrospinal-fluid pressure, possibly reflecting intermittent pressure surges. It then states the gait-cognition-urinary triad and immediately qualifies its order: gait apraxia is generally initial, consistent and prominent, whereas cognitive decline clearly preceding gait dysfunction makes the diagnosis less likely.

The gait description earns clinical depth. Mention impaired alternation, failure of forward weight transfer, repetition of movement with the same leg, a stuck weight-bearing foot and a multi-step turn, while stepping over a stick remains possible. Characterise cognition as slowed and subcortical with relative language preservation and common abulia. Describe urinary urgency and frequency progressing to incontinence.

Complete the answer with investigation and management logic. Imaging demonstrates ventricular enlargement, especially at the temporal horns, but cannot diagnose the disorder alone because the appearance is nonspecific and may resemble hydrocephalus ex vacuo. Cerebrospinal-fluid pressure, protein and glucose are normal. Large-volume removal or serial lumbar punctures are used as a predictive test: gait improvement supports the diagnosis and predicts shunt benefit, lack of cognitive improvement does not negate a positive gait response, and a negative test does not exclude the disorder. Permanent ventricular drainage may relieve the syndrome, but benefit is not fully predictable and complications can be serious.

> In normal-pressure hydrocephalus, gait leads the diagnosis and usually the response; imaging supports, temporary drainage predicts, and a negative tap test does not exclude.

17 · Rapidly progressive and prion dementias, and the genetic subcortical dementias

Why the pace changes the whole problem. A dementing illness that accelerates over a short interval is not simply a common dementia running faster. It changes the prior probabilities, raises the cost of delay and demands an intensive search for inflammatory, neoplastic, autoimmune, nutritional, infectious, vascular and prion causes. **Creutzfeldt-Jakob disease** (CJD) must remain one possibility among several throughout that urgent search.

The diagnostic task has two simultaneous aims. The first is rescue: identify processes in which prompt causal treatment may alter outcome. The second is recognition: when rapidly progressive cognitive failure is accompanied by multifocal neurological signs and a characteristic biomarker pattern, investigate prion disease directly. A strong answer therefore moves from tempo and syndrome to targeted exclusion, then to prion phenotype, imaging, cerebrospinal-fluid seeding assay and the level of diagnostic certainty.

PACE	PATTERN	REVERSIBLE	PRION
DEFINE THE PROBLEM	LOCALISE THE SYNDROME	EXCLUDE EARLY	TEST DIRECTLY

17.1 Rapidly progressive dementia is a diagnostic emergency

Rapidly progressive dementia (RPD) is a clinical problem defined primarily by its relatively abrupt onset and progression, not by a single aetiology. Its differential is distinct from that of a typically insidious dementia, and the patient requires intensive investigation. A **6-12 months** development is a red flag against a straightforward diagnosis of a conventional slowly progressive dementia. At the still faster end, multifocal neurological deficits emerging over days to weeks strongly raise suspicion of sporadic CJD.

Tempo must be reconstructed rather than guessed from the referral label. Establish the previous cognitive and functional baseline, the first unequivocal change, the order in which cognitive, behavioural, gait, visual, cerebellar and motor abnormalities appeared, and whether decline has been smooth, stepwise or punctuated. Apparent rapidity may reflect late recognition of a longer illness, but true acceleration must be taken seriously. A short downhill course is a useful warning sign for CJD. Tempo alone cannot establish the diagnosis.

RULE **In RPD, search for a treatable cause and investigate prion disease in parallel.** Prion disease should be considered whenever dementia or ataxia progresses rapidly and no better explanation has emerged; serial exclusion alone must not postpone the tests that can support an accurate antemortem diagnosis.

Prion disease can overlap clinically with other neurodegenerative syndromes, including Alzheimer disease, dementia with Lewy bodies and frontotemporal dementia. This overlap is precisely why pace alone is insufficient. The clinician must ask whether the illness has acquired the multifocal neurological texture expected of prion disease and whether the biomarker pattern converges. Conversely, a familiar dementia phenotype should not reassure when the trajectory is implausibly fast.

17.2 The RPD differential: preserve breadth before narrowing

The differential should be organised by mechanism because an unstructured list is easy to truncate. The categories below are not equally likely in every patient, but each represents a route by which a rapidly dementing presentation can arise. Clinical context then determines the order and intensity of testing.

CATEGORY	CAUSES TO NAME	REASONING EMPHASIS
Inflammatory	Cerebral vasculitis, multiple sclerosis, sarcoidosis	Seek evidence that cognition belongs to a broader inflammatory neurological syndrome.
Neoplastic	Primary central nervous system tumour, cerebral metastases, paraneoplastic limbic encephalitis	Consider direct structural disease and remote tumour-associated neurological illness.
Autoimmune	Voltage-gated potassium-channel antibody syndrome	Psychiatric or cognitive prominence can accompany an immune-mediated illness.
Nutritional	Thiamine deficiency and Wernicke-Korsakoff syndrome	Name the deficiency while keeping its disease teaching with the relevant owner chapters.
Infectious	Cerebral abscess, herpes encephalitis, progressive multifocal leucoencephalopathy, HIV, subacute sclerosing panencephalitis, Whipple disease	Use exposure, host context and accompanying neurological features to direct exclusion.
Prion	Creutzfeldt-Jakob disease	Look for rapid multifocal decline and convergent MRI, seeding-assay and other biomarker evidence.
Vascular	Multiple infarcts, including embolic disease related to endocarditis, and CADASIL	Interrogate the temporal pattern and the possibility of recurrent vascular injury.

This is a triage framework, not permission to teach neighbouring disorders inside an RPD answer. The autoimmune and limbic encephalitis workup belongs in the Neuropsychiatry of systemic and neurological disease notes. The endocrine, metabolic and nutritional diseases themselves belong in the Endocrine and metabolic psychiatry notes; the Korsakoff amnestic syndrome and the vascular neurological syndromes belong in the Stroke, TBI, delirium and focal brain syndromes notes. Here their role is to remain visible as competing, potentially actionable explanations.

- First decide whether the pace is genuinely rapid relative to the patient's prior baseline.
- Next identify the dominant accompanying syndrome: encephalitic, cerebellar, multifocal motor, visual, vascular or systemic.
- Then investigate treatable categories while obtaining prion-directed studies when the clinical pattern warrants them.

PITFALL

Do not wait for every reversible investigation to return before requesting diffusion-weighted MRI and a prion seeding assay in a convincing RPD. Sequential testing wastes the information contained in tempo; parallel testing preserves rescue and recognition.

17.3 Human prion disease: proportions and clinical frame

Human prion diseases are rare across populations, occurring at a frequency of 1 to 1.5 cases per million; the United States records approximately 300 diagnosed cases annually. Most are sporadic. The triangulated transmission breakdown is **about 85-90% sporadic, 14-15% genetic and 1% acquired**. These proportions matter because absence of family history or recognised exposure does not meaningfully reduce suspicion for the dominant sporadic form.

CJD may be classified as **sporadic CJD** of unknown aetiology, autosomal-dominant familial CJD, iatrogenic CJD related to contaminated biological material or instruments, and variant CJD associated with consumption of cattle contaminated by bovine spongiform encephalopathy. The transmission label predicts more than provenance. It influences age profile, presenting syndrome, disease duration and the yield of familiar investigations.

Anatomically, CJD affects the basal ganglia and thalamus, the neocortical mantle with frontal and posterior cortical involvement, and the cerebellum. This distributed involvement explains why a single patient may show cognitive deficits, neuropsychiatric change, altered sleep, visual disturbance and gait instability. The syndrome therefore combines cortical and subcortical dysfunction, often with cerebellar participation; reducing it to “dementia plus jerks” loses its diagnostic breadth.

MNEMONIC

PACE: Progression, Associated signs, Convergent tests, Exposure and inheritance

Progression identifies RPD; **A**ssociated signs reveal the multifocal syndrome; **C**onvergent tests support prion disease; **E**xposure and inheritance distinguish acquired and genetic routes. This is an examination spine, not a diagnostic instrument.

17.4 Sporadic CJD: the evolving multifocal syndrome

Sporadic CJD usually begins insidiously with nonspecific behavioural or constitutional change, including anxiety, depression, altered sleep, reduced appetite, fatigue, dizziness, weight loss and social regression. Forgetfulness and failure of higher cortical functions then become clearer. Approximately a third have prodromal headache, fatigue or weight loss, while almost a third initially show purely neurological deficits, especially cerebellar gait ataxia. Thus an early psychiatric or cerebellar presentation does not contradict the later diagnosis.

Onset is usually within **45-90 years, peaking around 65 years**. The classic established picture is rapidly progressive multifocal dementia with myoclonus and gait ataxia. Additional visual, cerebellar, pyramidal and extrapyramidal signs accumulate as anatomical involvement widens.

Startle myoclonus, precipitated by loud noise or touch, is prominent but is often absent early and emerges with progression. Epileptic seizures are uncommon, occurring in only 10%.

Uncommon clinical forms are worth recognising without mistaking them for separate aetiologies. The Heidenhain variant has prominent occipital disease with cortical blindness, visual hallucinations and agnosia. The Brownell-Oppenheimer variant presents as a pure cerebellar syndrome associated with widespread cerebellar pathology. Their examination value is simple: focal prominence early in the course does not exclude a disease that later becomes multifocal.

The downhill trajectory is severe. Progression over weeks may culminate in akinetic mutism and death, often within 2-3 months. Around **70% die in under 6 months and 90% by 1 year**, although duration is variable and a minority have a much longer course. The pace helps distinguish the syndrome clinically, but variability prevents duration from functioning as an absolute exclusion criterion.

17.5 Diagnostic criteria and levels of certainty

The clinical-sign category begins with rapidly progressive dementia and asks for myoclonus, cerebellar or visual signs, pyramidal or extrapyramidal signs, and akinetic mutism. The test category includes periodic sharp-wave complexes on EEG, cerebrospinal-fluid **14-3-3 protein**, characteristic diffusion-weighted or FLAIR MRI abnormalities, and a positive real-time quaking-induced conversion assay.

DSM-5-TR requires motor features of prion disease, namely myoclonus or ataxia, or biomarker evidence when diagnosing major or mild neurocognitive disorder due to prion disease. This diagnostic-manual formulation identifies the necessary neurological or biological support. The more operational sporadic-CJD case definitions below show how clinical signs, tests, duration and seeding evidence are combined into levels of certainty.

DIAGNOSTIC LEVEL	REQUIRED CONSTRUCTION	INTERPRETIVE POINT
Probable sporadic CJD, clinical plus tests	Rapidly progressive dementia, at least 2 clinical signs and at least 1 qualifying test	The diagnosis rests on convergence, not on one nonspecific marker.
Probable sporadic CJD, seeding route	A progressive neurological syndrome with positive cerebrospinal-fluid RT-QuIC	A positive seeding assay can establish the probable category through the alternative route.
Possible sporadic CJD	Rapidly progressive dementia, at least 2 clinical signs and duration shorter than 2 years	This level lacks the qualifying test evidence required by the clinical-plus-tests route.
Definite prion disease	Characteristic neuropathology with prion-protein immunohistochemistry, or a confirmed PRNP mutation in inherited disease	Definitive classification remains tissue- or mutation-based.

Neuropathological confirmation requires classical spongiform change, neuronal loss, astrocytosis and positive prion-protein immunohistochemistry on biopsy or autopsy. Yet modern imaging,

seeding assays and PRNP genotyping permit accurate antemortem diagnosis in the large majority, so brain biopsy is now rarely required. Probable CJD is confirmed as definite at autopsy in more than 95% of cases. The distinction is not between a useless antemortem label and certainty after death; it is between high clinical confidence and the formal requirements for definite disease.

17.6 MRI, cerebrospinal fluid and EEG: use convergence, not hierarchy alone

Brain MRI with diffusion-weighted sequences is the **most sensitive non-invasive diagnostic tool** for sporadic CJD. Typical findings are restricted diffusion in the basal ganglia and thalami with cortical ribboning; regional combinations differ between patients. Signal hyperintensity may also occur in the same regions on T2-weighted or FLAIR images, but is less striking. Imaging modality physics belongs in the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes; the dementia-specific point is the distribution and diagnostic contribution of the abnormal signal.

Routine cerebrospinal-fluid cell count and biochemistry remain normal in sporadic CJD. This does not make lumbar puncture unhelpful: routine examination is essential for excluding other causes of rapid cognitive decline, while disease-associated proteins and seeding assays address a different question. The procedure and generic study of cerebrospinal fluid belong in the Neurology foundations for psychiatrists notes. Routine blood tests are usually normal and, when abnormal, have no particular diagnostic value for sporadic CJD.

Real-time quaking-induced conversion (RT-QuIC) is an amyloid-seeding assay. Repeated shaking and incubation amplify prion seeds from patient cerebrospinal fluid using recombinant prion protein as substrate and a fluorescent thioflavin-T readout. Its sensitivity is comparable to a positive cerebrospinal-fluid 14-3-3 result, but its specificity for sporadic CJD approaches **100%**. Nasal-brushing and skin-seeded assays also show performance approaching that level, but are not yet routine clinical tests.

Protein 14-3-3 and tau are commonly used cerebrospinal-fluid biomarkers with high sensitivity but variable specificity. EEG may show periodic, often triphasic synchronous discharges at 0.5-2 Hz at some point in the illness. These findings must be interpreted as components of a syndrome. A nonspecific injury marker cannot carry the diagnosis alone, and the expected marker pattern differs across sporadic, genetic and variant disease.

TRAP **Do not memorise a single textbook percentage for periodic EEG activity as settled fact.** One standard source reports pseudoperiodic sharp waves in a minority of sporadic cases, whereas another reports the finding substantially more often late in illness. The verified diagnostic-manual statement is qualitative: periodic triphasic synchronous discharges may appear during the course. Quote the pattern and frequency band, not a falsely reconciled prevalence.

GOOD TO REMEMBER

A positive RT-QuIC has far greater rule-in value than a nonspecific rise in an injury-associated cerebrospinal-fluid protein. Still, the best answer is convergent: phenotype, diffusion-weighted MRI and seeding evidence interpreted together.

17.7 Acquired disease: route changes the phenotype

Iatrogenic CJD should be considered in a progressive, predominantly cerebellar syndrome with behavioural disturbance after exposure to human pituitary hormones, human dura mater, an implicated corneal graft or neurosurgical instruments used on a definite or probable prion case. Limb dysaesthesiae are frequent. Incubation may exceed 40 years, so a remote exposure remains relevant.

Route explains the clinical contrast. Intracerebral or optic inoculation produces a classical rapidly progressive dementing picture and may have incubation shorter than 2 years. Peripheral inoculation, especially through pituitary-derived growth hormone, more often produces a progressive cerebellar syndrome reminiscent of kuru, with incubation typically 15 years or longer. MRI may resemble sporadic CJD, but periodic sharp waves are rarely seen. Cerebrospinal-fluid 14-3-3 is detected in only approximately half, and RT-QuIC sensitivity is substantially reduced compared with sporadic disease.

Variant CJD is clinically different. Depression, anxiety, withdrawal or delusions typically precede dementia, and painful sensory symptoms, cerebellar ataxia and abnormal movements may also lead the cognitive syndrome. It predominantly affects young adults and follows a course longer than 6 months. Average onset is 29 years, compared with 65 years in sporadic CJD; neurological features usually follow the initial psychiatric presentation within 6 months, and median survival is 13 months.

The EEG shows nonspecific slowing rather than the periodic sharp-wave pattern of sporadic disease, and cerebrospinal-fluid 14-3-3 may be raised or normal. FLAIR MRI may show the **pulvinar sign**, with thalamic signal more prominent than basal-ganglia signal, in approximately 90%; it may appear late and is not pathognomonic. Tonsillar tissue may demonstrate characteristic prion-protein findings. Unlike sporadic CJD, variant disease can be transmitted by blood transfusion.

17.8 Genetic disease, the Huntington boundary and care priorities

Genetic prion disease accounts for roughly the inherited share already described and may be established by a confirmed **PRNP mutation**. Up to 15% of prion disease cases carry autosomal-dominant mutations in this gene. Family history therefore matters, but its absence should not displace the much larger sporadic category. Genetic counselling detail and mutation-specific syndromes cannot be supplied from the present evidence base and should not be reconstructed from memory.

Huntington disease belongs among the common causes of dementia and within the basal-ganglia-predominant neurodegenerative dementias. In this section it is a pointer, not an

invitation to repeat its genetics, motor-psychiatric-cognitive triad or staging. Full teaching belongs in the Neuropsychiatry of systemic and neurological disease notes. The dementia-lens point is that a genetic subcortical neurodegenerative disorder sits beside prion disease in the broader differential, but its usual disease teaching must remain separate.

Care follows diagnostic clarity. LOCUS depends on clinical instability and diagnostic needs: urgent hospital assessment is appropriate when rapidly progressive cognitive failure and neurological deficits require intensive investigation, while later care may move across specialist, community and palliative settings. FOCUS is to identify a treatable mimic early, establish the level of prion-disease certainty, address distressing cognitive, behavioural and neurological disability, communicate prognosis honestly and support those providing care. MODUS is diagnostic, supportive, psychological and social; no disease-modifying treatment currently exists for CJD.

- Explain that a probable antemortem diagnosis can be highly accurate even though definite classification remains neuropathological or mutation-based.
- Use phenotype, MRI and RT-QuIC together, and interpret negative or low-yield markers in light of the suspected prion subtype.
- Keep reversible RPD causes active until reasonably excluded, while avoiding futile delay once convergent prion evidence is present.

> Rapid dementia demands parallel reasoning: exclude treatable causes urgently, but investigate prion disease directly when rapid multifocal decline, diffusion restriction and prion-seeding evidence converge.

The compact examination answer is therefore coherent rather than encyclopaedic. Define RPD by its abnormal pace, organise the broad differential, describe the evolving multifocal syndrome of sporadic CJD, state the probable and definite diagnostic routes, compare the information supplied by MRI, routine cerebrospinal fluid, 14-3-3, RT-QuIC and EEG, and then distinguish iatrogenic and variant patterns. End by recognising inherited prion disease and naming Huntington disease only at the dementia boundary. That sequence demonstrates the urgency and restraint this topic tests.

18 · Mild cognitive impairment, subjective cognitive decline and the cognitive continuum

Why it matters clinically. The borderline between expected cognitive experience and dementia is not a single test score. Patients may notice decline before testing confirms it; measurable impairment may later appear while everyday independence remains intact; some then progress, some remain stable, and some no longer meet the earlier label. The marks lie in defining each state precisely, separating subjective report from objective performance, and showing why function and trajectory matter more than an isolated complaint.

The useful model is a cognitive continuum, but not an inevitable conveyor belt. **Subjective cognitive decline** describes a perceived worsening without an objective deficit on

neurocognitive testing in the descriptive evidence available here. **Mild cognitive impairment** adds objective or longitudinally demonstrated cognitive decline while basic everyday activities remain preserved and the disturbance does not warrant dementia. **DSM-5-TR mild neurocognitive disorder** occupies the corresponding diagnostic territory, with modest decline but preserved independence. Dementia lies beyond the functional boundary and is discussed elsewhere in this chapter. These labels organise uncertainty; they do not by themselves disclose aetiology or destiny.

Concern	Testing	Function	Course
SUBJECTIVE CHANGE	OBJECTIVE CHANGE	INDEPENDENCE	TRAJECTORY

18.1 The continuum is multidimensional, not merely ordinal

A linear diagram is useful for orientation, but clinical reasoning requires several dimensions. Subjective experience asks whether the person perceives decline from a previous level. Objective assessment asks whether performance is impaired or has demonstrably worsened over time. Functional assessment asks whether the person remains independent, perhaps through greater effort or compensatory strategies. Longitudinal assessment asks whether the state progresses, persists, fluctuates or resolves. A patient can move differently along each axis.

The distinction between impairment and disability is central. Poor performance on a cognitive task establishes neither dementia nor loss of independence. Conversely, a person may report meaningful decline while remaining within expected limits on the tests administered. The clinician must then decide whether the complaint reflects a very early cognitive change, a limitation of the assessment, ordinary variation, or another explanatory condition. The correct response is a formulation with explicit uncertainty, not a forced categorical answer.

RULE **Place the patient on the continuum using complaint, objective cognition, everyday independence and longitudinal course together.** No single component can substitute for the others, and the label should change when later evidence changes the formulation.

This also explains why “pre-dementia” can mislead. It is acceptable as a risk-oriented shorthand only if it does not imply that every patient will progress. Mild impairment is clinically important because the probability of future dementia is higher in several studied cohorts, yet community observations include stability and reversion. The continuum is therefore a map of current cognitive-functional state and future risk, not a declaration of a predetermined biological sequence.

18.2 Subjective cognitive decline: symptom before measurable deficit

In the cited textbook account, subjective cognitive impairment, also termed subjective cognitive decline, is described as the earliest symptomatic manifestation before mild cognitive impairment in a proposed Alzheimer disease sequence. Its defining contrast is straightforward: the person reports worsening memory, but neurocognitive testing does not demonstrate a deficit. This staging statement arose in a biomarker-research discussion and should not be inflated into proof that every such complaint represents that disease.

Subjective memory complaints are common among community-dwelling older adults. Typical concerns involve remembering names, finding the intended word, locating objects and sustaining concentration compared with earlier functioning. The phenomenology is familiar, but familiarity does not make it trivial. The clinician should establish what has changed, when the change was first noticed, whether others observe it, which situations expose it, and whether the person has altered routines to compensate.

Terminology is inconsistent. Some experts use mild cognitive impairment broadly for subjective complaints occurring without overt dementia or obvious affective, anxiety, sleep, pain or attention-related explanations. Others reserve the term for an objectively impaired state or for an early neurodegenerative syndrome. This is not pedantry. Different entry definitions assemble different populations, and those populations will produce different apparent progression rates. A candidate should therefore define the term before quoting prognosis.

PITFALL

Do not tell a patient that a subjective complaint with normal testing is "early Alzheimer disease." The permitted evidence supports a descriptive pre-MCI research stage, not an individual aetiological diagnosis. First describe the complaint, test objectively, assess alternative explanations and observe the trajectory.

The available evidence does not provide a formal operational feature set for subjective cognitive decline. Accordingly, a list of remembered "plus" features should not be manufactured. The defensible definition remains descriptive: perceived decline from prior ability without objective impairment on the assessment used. The limitations of that assessment and the need for follow-up should be stated openly.

18.3 Defining mild cognitive impairment and mild neurocognitive disorder

The **Report of the International Working Group on MCI** requires cognitive decline that does not warrant a dementia diagnosis, or evidence of decline over time on objective cognitive tasks, while basic activities of daily living remain preserved. This formulation is deliberately broader than a memory-only syndrome. It makes two points that survive changes in terminology: decline must be demonstrated clinically or objectively, and basic function remains sufficiently preserved to keep the syndrome on the mild side of the dementia boundary.

The diagnostic-manual formulation requires modest decline from prior performance in at least one cognitive domain. Evidence must come from concern expressed by the patient, an informant or the clinician and from modest impairment on standardised testing or another quantified clinical assessment. The domains considered are complex attention, executive function, learning and memory, language, perceptual-motor function and social cognition. Deficits do not interfere with independence in everyday activities, although greater effort, accommodation or compensation may be necessary. The syndrome must not occur exclusively during delirium and must not be better explained by another mental disorder. Aetiology is then specified rather than assumed.

The older **Mayo Clinic Alzheimer's Disease Research Center criteria** were narrower and memory-centred. They required a memory complaint, preferably corroborated; objective memory impairment adjusted for age and education; preserved general cognitive function; intact everyday activities; and absence of dementia. Their historical importance is conceptual: they identify the classic amnestic presentation from which the broader domain-based framework developed.

AXIS	INTERNATIONAL WORKING-GROUP FORMULATION	DIAGNOSTIC-MANUAL MILD DISORDER	ORIGINAL MEMORY-CENTRED CRITERIA
Core change	Cognitive decline insufficient for dementia, or objective decline over time	Modest decline from prior performance	Memory complaint with objective memory impairment
Evidence	Clinical decline or serial objective decline	Concern plus standardised testing or quantified clinical assessment	Complaint preferably corroborated, with age- and education-adjusted impairment
Cognitive scope	Not restricted to memory	Attention, executive, learning-memory, language, perceptual-motor or social-cognitive domains	Memory-centred, with general cognition preserved
Functional boundary	Basic everyday activities preserved	Independence preserved despite possible extra effort or compensation	Everyday activities intact
Exclusions	Does not warrant dementia	Not exclusively delirium and not better explained by another mental disorder	Not demented

The apparent differences are best understood as evolution, not contradiction. The early criteria selected an amnestic group; later formulations recognise impairment outside memory and place more explicit weight on quantified assessment and maintained independence. In an answer, state which framework is being used rather than mixing elements into an unnamed hybrid.

18.4 Function is the practical boundary

Why function outranks the raw score. The mild diagnosis is not defined by a universal cognitive cutoff in the evidence supplied here. Its practical boundary is independence. A person may take longer, rely on written prompts, reorganise tasks, avoid cognitively demanding situations or need greater effort while still completing essential activities independently. Such compensation is diagnostically informative: it can reveal decline without automatically converting mild impairment into dementia.

Functional questioning should be concrete and anchored to the previous level. Ask how the patient manages complex routines, new information, appointments, travel, communication, household organisation and occupational demands. The issue is not whether another person

sometimes helps, but why help is needed and whether cognitive deficits now interfere with independent performance. Physical limitation, sensory impairment, low mood, reduced opportunity and family overprotection can all change activity without proving cognitive dependence.

Objective cognition must likewise be interpreted as change from prior ability. Education, language, cultural familiarity, sensory status and test conditions influence performance. General administration and scoring of cognitive screens belongs in the Psychometry, Assessment and Descriptive Psychopathology notes. Here, the relevant point is narrower: testing corroborates and characterises decline; it does not replace history, functional analysis or serial observation.

GOOD TO REMEMBER

Document compensation, not merely success or failure. "Independent with a diary and repeated checking" carries more information than "activities intact." It captures the diagnostic-manual distinction between preserved independence and the greater effort used to maintain it.

Exclusion is equally important. A mild cognitive syndrome should not be diagnosed when the disturbance occurs only in delirium or is better explained by another mental disorder. The full comparison with delirium and depression is taught elsewhere in this chapter, while delirium itself belongs to the Stroke, TBI, delirium and focal brain syndromes notes. The aim here is not to teach those disorders, but to prevent their presence from being erased by a convenient cognitive label.

18.5 Phenotyping: amnestic, nonamnestic and domain burden

Once mild impairment is established, phenotype it rather than stopping at the umbrella label. **Amnestic MCI** is dominated by memory impairment, whereas **nonamnestic MCI** is dominated by nonmemory impairment such as executive or language dysfunction. Each may show a **single-domain or multiple-domain pattern**. This classification describes the pattern of impairment and can guide the aetiological hypothesis, but it is not itself an aetiological diagnosis.

The classic mapping links single-domain amnestic impairment more closely with underlying Alzheimer disease pathology. Amnestic impairment affecting multiple domains may be associated with that pathology or with vascular dementia. Nonamnestic impairment is related more closely to dementia with Lewy bodies or frontotemporal neurodegeneration. These are associations, not conversion rules. Detailed criteria for each dementia remain with their disorder-specific sections.

The domain distinction matters because "memory clinic" language can obscure nonmemory presentations. Executive inefficiency may first appear as difficulty planning or shifting strategy. Language decline may appear as impaired retrieval or expression. A person can therefore satisfy a mild cognitive framework without the classic memory-centred picture. Conversely, a memory complaint alone does not establish amnestic impairment until objective assessment confirms the relevant deficit.

MNEMONIC

COFI: Complaint, Objective change, Function, Interval course

Complaint identifies the perceived decline; **O**bjective change establishes the cognitive phenotype; **F**unction locates the mild-versus-dementia boundary; **I**nterval course shows whether the label progresses, persists or recedes.

A robust formulation therefore has two layers. The syndromic layer states that mild cognitive impairment or mild neurocognitive disorder is present and specifies the domains affected. The aetiological layer states what process is suspected and how certain that suspicion is. Keeping the layers separate prevents a common shortcut in which an amnestic phenotype is silently rewritten as confirmed neurodegenerative disease.

18.6 Prognosis: progression, persistence and reversion

Prognosis is probabilistic and population-dependent. In the Mayo Clinic referral sample, conversion to dementia was **12% per year compared with 1% to 2% per year in controls**. A separate account in Kaufman's gives **10% per year compared with 1% to 2% per year in unaffected peers**. These figures are similar in direction, but they are not interchangeable estimates from the same population.

Community findings make the uncertainty more visible. In the Monongahela Valley Independent Elderly Survey, **10% to 17% progressed to dementia, while 33% to 56% became non-MCI**. Thus reversion is not a theoretical footnote. It may reflect genuine improvement, variable classification, fluctuating contributors or regression around a threshold. Whatever the explanation, it invalidates language of inevitability.

Stahl's estimates that **15% to 20% of people aged 65 years or older have MCI**. On clinical grounds without biomarkers, the same source gives **6% to 15% conversion to dementia each year**, with additional cumulative estimates that differ according to follow-up and outcome definition. The wide spread is understandable when some authors include subjective complaints while others require objective impairment or select specialty referrals.

TRAP **Do not quote one MCI conversion rate as universally correct.** Kaplan and Sadock report a referral-sample estimate, Kaufman's gives a separate point estimate, and Stahl's gives a broader clinically defined range without biomarkers. Name the source, setting, case definition, follow-up frame and outcome before using the figure.

EVIDENCE CONTEXT	WHAT THE ESTIMATE SHOWS	WHY IT MUST NOT BE UNIVERSALISED
Specialty referral sample	Higher annual conversion than controls	Clinical selection enriches risk and differs from community ascertainment
Textbook point estimate	Higher annual progression than unaffected peers	A summary estimate does not erase study and definition differences
Community survey	Progression coexists with loss of the MCI label and return to normal	Community cases are broader and the state may not remain stable
Clinical definition without biomarkers	A range of annual conversion and substantial cumulative risk	Subjective and objective case definitions may be mixed across accounts

For the individual patient, predictors and biomarkers may refine risk, but the permitted evidence does not justify a personalised numerical forecast. The defensible communication is that risk is elevated, outcomes vary, and reassessment matters. A change in label over time is new clinical information, not proof that the earlier clinician necessarily made an error.

18.7 Mild behavioural impairment broadens the prodromal lens

Cognitive change is not the only possible early signal. **Mild behavioural impairment** describes the emergence of neuropsychiatric symptoms that can precede cognitive decline and eventually dementia. It can accompany the cognitive deficits of mild cognitive impairment or draw attention before those deficits are obvious. The construct encourages clinicians to regard new behavioural change as potentially informative rather than automatically dismissing it as personality.

The operational framework groups symptoms into motivation, affect, impulse control, social appropriateness, and perception or thought content. A dedicated checklist is available for rating this construct. These categories are not a substitute for psychiatric diagnosis, and the presence of a symptom does not prove neurodegeneration. Their value lies in describing a change from longstanding behaviour, establishing its context, and integrating the behavioural trajectory with cognitive and functional findings.

Mild behavioural impairment also guards against a memory-only model of the continuum. A patient may first present because of apathy, emotional change, disinhibition, socially inappropriate conduct or unusual perceptual and thought experiences. The cognitive examination may then reveal an amnestic or nonamnestic phenotype, remain objectively normal, or change later. The behavioural construct and cognitive labels answer different questions and may coexist.

- Establish that the behaviour is new rather than lifelong.
- Describe the behavioural domain and its consequences without inferring a dementia subtype from it alone.

- Assess cognition and everyday independence in parallel, then review the behavioural and cognitive trajectories together.

The broader lesson is that early neurocognitive illness can be cognitively or behaviourally signalled, but neither signal is self-interpreting. Syndromic description should precede aetiological attribution. This protects against both under-recognition and premature diagnostic certainty.

18.8 Constructing and communicating the final formulation

A good formulation begins with the current state, not the feared destination. State whether there is a subjective complaint, objective decline, preserved independence and a demonstrable longitudinal change. If objective impairment is present, name the affected domains and whether the pattern is amnestic or nonamnestic, single-domain or multiple-domain. Then state exclusions and the degree of confidence in any suspected aetiology.

The history and assessment should answer a compact sequence:

- What has changed from the previous baseline, and is that change subjective, objectively demonstrated or documented serially?
- Does cognitive difficulty interfere with independence, or is independence maintained through effort and compensation?
- Which course is emerging: progression, persistence, fluctuation or reversion?

Communication should separate present diagnosis from future risk. “Your testing shows mild impairment, but you remain independent” is a current-state statement. “Some people progress while others remain stable or improve” is a prognosis statement. “The pattern may suggest a particular cause, but does not prove it” is an aetiological statement. Keeping these sentences separate reduces fatalism and false reassurance.

The same discipline strengthens an examination answer. Define the subjective and objective states, identify preserved independence as the boundary, describe the amnestic and nonamnestic phenotypes, and then discuss prognosis with source and population attached. Add behavioural change as a parallel early signal. Do not drift into detailed dementia biomarkers, imaging, drug treatment or disorder-specific diagnostic criteria, which belong to their dedicated sections.

> MCI means objective or longitudinal cognitive decline without loss of everyday independence; it raises dementia risk but does not make progression inevitable.

The cognitive continuum is therefore best understood as a revisable clinical model. Subjective concern can precede measurable impairment; measurable impairment can coexist with independence; behavioural change can accompany either; and later observation may show progression, stability or reversion. Precision at the borderline does not require false certainty. It requires explicit definitions, careful functional analysis and respect for longitudinal evidence.

19 · Delirium versus dementia versus depression: comparative differentiation

The marks lie in choosing the right discriminator, not in reciting three definitions. Delirium, dementia and depression can all present with poor recall, reduced participation, behavioural change and apparent functional decline. The useful comparison therefore begins with the architecture of the syndrome: consciousness, attention, speed of onset, fluctuation, progression, testing context and what happens over time. A candidate who starts and ends with "memory impairment" has selected the least discriminating common feature.

This section owns the distinction among the three syndromic presentations. It does not reproduce the full diagnosis, investigation or management of delirium, which belongs in the Stroke, TBI, delirium and focal brain syndromes notes. It also does not reteach depressive disorders or general cognitive-screen administration. Its purpose is narrower and clinically more useful: identify which observations change the probability of each explanation, recognise coexistence, and avoid converting a bedside clue into an absolute rule.

19.1 Begin with consciousness and attention

The first division is between cognitive impairment occurring with altered consciousness and cognitive impairment occurring in clear consciousness. In **delirium**, cognitive change accompanies an altered level of consciousness, alertness fluctuates, onset is generally sudden and acute, and **attention** is the hallmark impaired cognitive domain. In **dementia**, cognition is impaired in clear consciousness. This is the highest-yield conceptual distinction because a patient who cannot maintain or direct attention cannot be assessed as though a stable cognitive baseline were being sampled.

"Clear consciousness" does not mean that every person with dementia is calm, fully oriented or free of behavioural symptoms. It means that altered consciousness is not part of dementia by definition. Conversely, altered consciousness is not merely severe forgetfulness. It is a change in the patient's level or stability of awareness and alertness accompanying cognitive impairment. The bedside sequence therefore runs from arousal to attention and only then to the pattern of memory, language and other cognitive findings.

RULE **Test the state before interpreting the score.** Altered consciousness, fluctuating alertness and impaired attention direct the formulation toward delirium; clear consciousness permits a dementia or depression comparison, but does not by itself settle it.

ATTENTION	CLEAR	ENCOURAGE	FOLLOW
DELIRIUM HALLMARK	DEMENTIA CONSCIOUSNESS	TESTING CONTEXT	COGNITIVE TRAJECTORY

19.2 Delirium and dementia: feature-by-feature comparison

The comparison below should be read as a pattern, not as a checklist in which every cell must be present. Its strongest use is relational: acute or subacute change plus fluctuation, altered alertness and impaired attention forms a different syndrome from insidious, slowly progressive decline in

otherwise clear consciousness. Memory impairment alone cannot make that distinction because episodic memory may be impaired in both.

CLINICAL AXIS	DELIRIUM PATTERN	DEMENTIA PATTERN
Onset	Acute or subacute	Insidious
Course	Fluctuating	Slowly progressive
Duration	Hours to weeks	Months to years
Alertness	Abnormally low or high	Typically normal
Sleep-wake cycle	Disrupted	Typically normal
Attention	Impaired	Relatively normal
Orientation	Impaired	Intact in early dementia
Working memory	Impaired	Intact in early dementia
Episodic memory	Impaired	Impaired
Thinking	Disorganised or delusional	Impoverished
Speech	Slow or rapid, and may be incoherent	Word-finding difficulty
Perception	Illusions or hallucinations are common	Typically unimpaired early; **Lewy body dementia** is an exception
Behaviour	Withdrawn or agitated	Varies with dementia type and is often intact early

Several rows deserve interpretation. Orientation and working memory can remain intact early in dementia, so preserved performance in these areas does not exclude a progressive disorder. Episodic memory can be impaired in both conditions, so asking only about delayed recall is diagnostically weak. Perceptual disturbance supports delirium in the broad comparison, but the exception of Lewy body dementia prevents a simplistic “hallucinations equal delirium” answer. Behaviour also overlaps: withdrawal and agitation may arise in delirium, while dementia-related behaviour depends on the underlying type and stage.

The table also explains why serial observation is often more revealing than a single polished interview. Fluctuation is not the same as ordinary variation in effort or mood. It belongs to a larger delirium pattern that includes altered alertness and impaired attention. Slow progression is likewise not established by one poor test; it is a longitudinal conclusion constructed from history, collateral information and change from baseline.

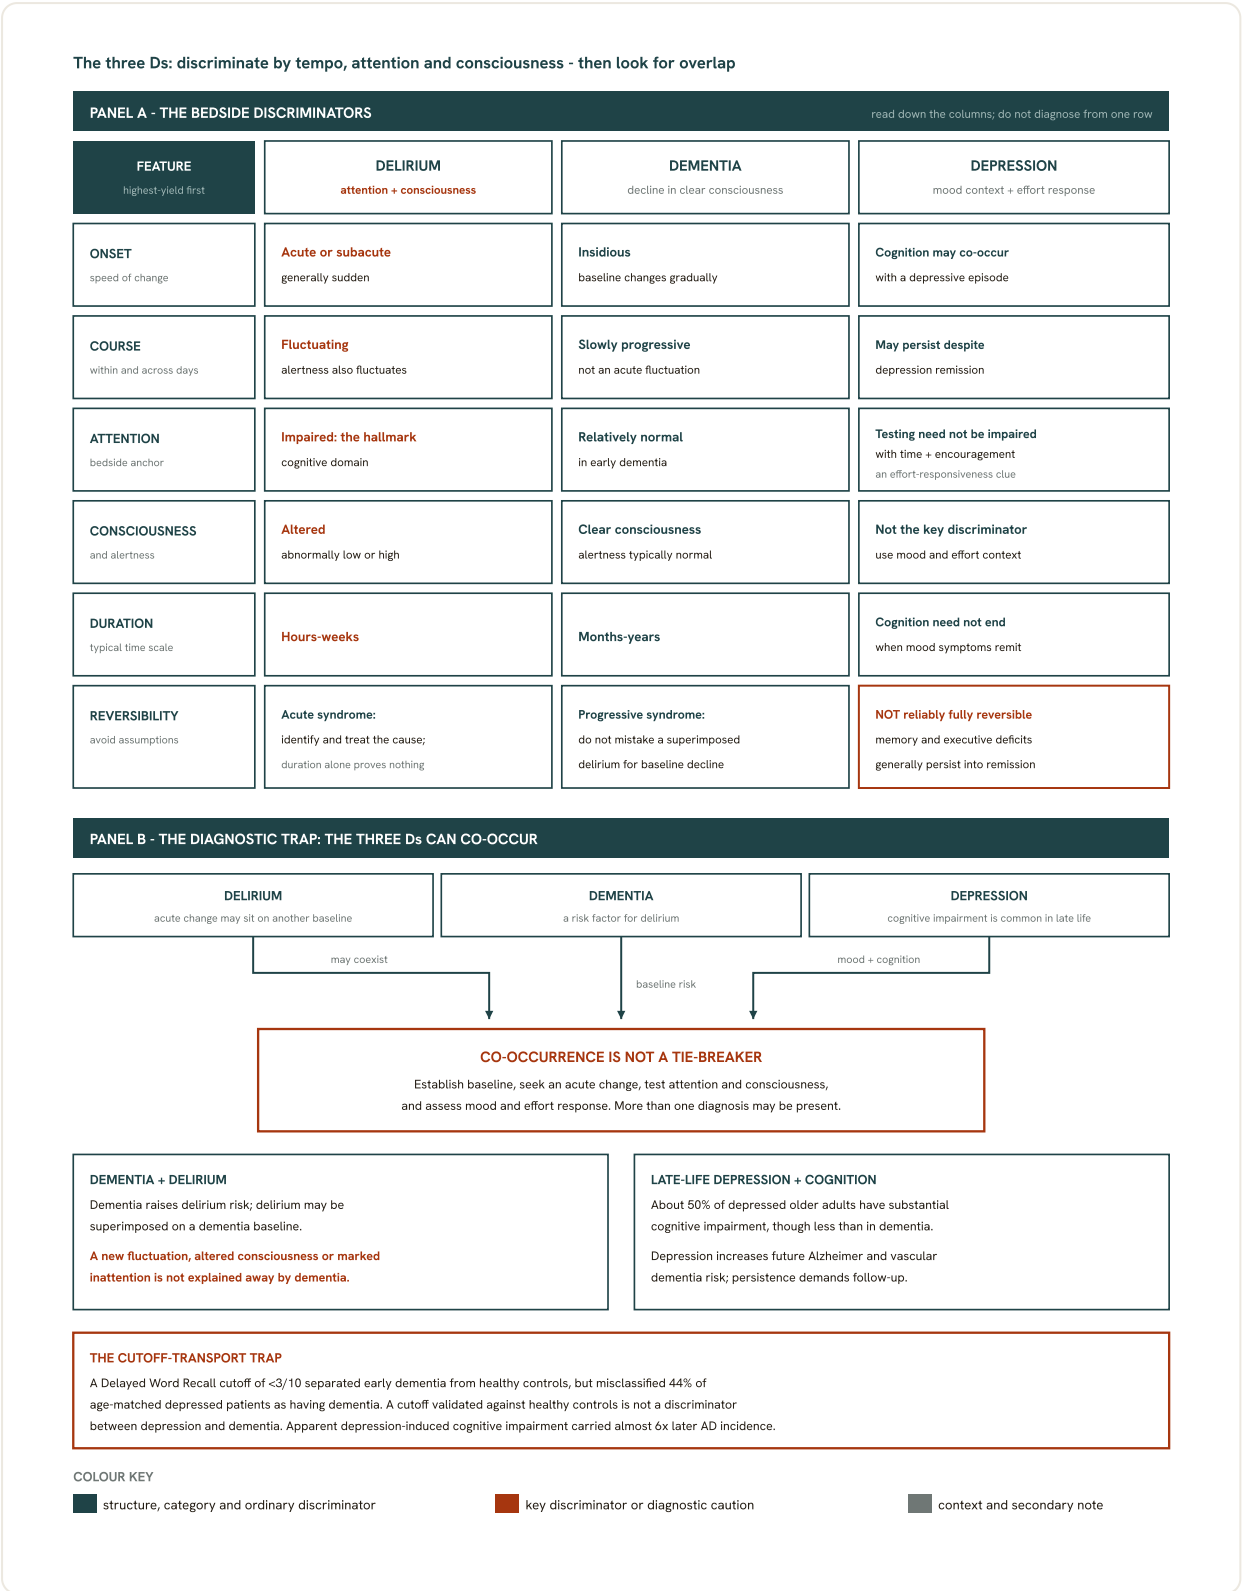


Fig 32.7 **The three Ds: delirium, dementia and depression.** Acute fluctuation, impaired attention and altered consciousness point to delirium; insidious progression in clear consciousness supports dementia. Depression-associated cognitive impairment may respond to time and encouragement during testing but is not reliably fully reversible, and overlap among the three conditions makes baseline and longitudinal follow-up essential. (Companion to Psychiatric Studies; Kaplan & Sadock's Comprehensive Textbook of Psychiatry 2024; Kaufman's Clinical Neurology for Psychiatrists.)

The accompanying figure should be used as a discriminator rather than as three sealed boxes. The decisive clinical move is to follow the branch that best explains the patient's state, while keeping open the possibility that more than one process is present.

19.3 The coexistence problem: delirium on a dementia baseline

Dementia is a risk factor for delirium, and delirium may be **superimposed on dementia**. This destroys a common false choice. A patient with established progressive cognitive decline can still develop a new acute, fluctuating disturbance of alertness and attention. The prior dementia explains vulnerability and baseline impairment; it does not explain away the new state change.

Clinically, the comparison must therefore be made against the patient's own baseline rather than against an imagined normal older adult. The question is not simply, "Is cognition impaired?" It is, "What has changed from the usual pattern, how rapidly, and in which domain?" A longstanding difficulty with memory may belong to the baseline dementia, while a new inability to sustain attention with unstable alertness indicates an additional syndrome. This reasoning is more dependable than treating the current examination as though it represented the patient's usual level.

Collateral history is valuable here because a fluctuating state may look deceptively coherent during a brief encounter. The informant's task is not to diagnose. It is to describe the baseline, the change and its time course. The examiner then maps that history onto the consciousness-attention-course framework. Full delirium assessment and management remain a separate topic in the Stroke, TBI, delirium and focal brain syndromes notes.

> **PITFALL**
>
> Do not attribute every acute behavioural or cognitive change in a person with dementia to "progression." Progression is characteristically slow, whereas a new fluctuating alteration in alertness and attention raises delirium superimposed on the established baseline.

19.4 Depression-associated cognitive impairment is not a disposable mimic

The older label **pseudodementia** implied that cognitive impairment accompanying depression was merely apparent and would fully reverse when mood improved. That assumption is unsafe. Substantial cognitive impairment is present in about **50%** of depressed older adults, although it is less marked than in dementia and occurs more frequently than in younger depressed patients. The problem is therefore common enough to be part of routine differential diagnosis rather than a rare examination curiosity.

More importantly, cognitive impairment accompanying late-life depression generally persists into remission, especially in memory and executive function. The term "pseudodementia" may still appear in clinical and examination language, but it should not carry the old promise of complete reversibility. Mood recovery and cognitive recovery are related observations, not interchangeable endpoints. A patient may improve affectively yet retain measurable cognitive difficulty.

This changes the interpretation of follow-up. Persistence after mood improvement does not automatically prove that dementia was present from the beginning, just as improvement with encouragement does not prove that the difficulty was wholly caused by depression. The source literature leaves unresolved whether depressed mood in this setting can represent prodromal

dementia or an etiologically related but distinct process. The appropriate stance is longitudinal and probabilistic: document both mood and cognition, observe their trajectories, and resist forcing either into the role of a perfect explanatory variable.

Depression also increases the future risk of both **Alzheimer disease dementia** and **vascular dementia**. This association is clinically important, but it does not permit retrospective certainty about cause in an individual patient. Depression may be a coexisting disorder, a risk marker, a related process or part of an evolving presentation. The distinction must therefore remain open to revision as the longitudinal pattern becomes clearer.

19.5 Bedside testing: effort-responsiveness is a clue, not a verdict

A useful bedside distinction is **effort-responsiveness**. Cognitive performance in depression may improve when the examiner provides sufficient time and encouragement. This contrasts with the more persistent test difficulty expected when an underlying dementia is driving the deficit. The clinical value lies in observing how performance changes under supportive testing conditions, not in judging the patient as motivated or unmotivated.

The examiner can use this clue without turning the encounter into an unstandardised contest:

- First establish whether alertness and attention are stable enough for the result to be interpretable.
- Then allow adequate time and offer neutral encouragement, noting whether engagement and performance improve.
- Finally interpret the response alongside history, functional change, mood and longitudinal course rather than as a stand-alone diagnostic sign.

The response to encouragement is not a licence to dismiss cognitive complaints. Late-life depression can be accompanied by substantial and persistent impairment, so the bedside clue must be held together with the longitudinal evidence. Nor should a poor performance under ordinary conditions be read as proof of dementia when the patient was slowed, discouraged or insufficiently engaged. The comparison is about the pattern of response, not a moral inference about effort.

> **GOOD TO REMEMBER**
>
> When depression and dementia are both plausible, record the testing conditions as carefully as the score. Adequate time, neutral encouragement and observed response can alter interpretation, while follow-up shows whether the cognitive problem tracks mood recovery or persists beyond it.

Detailed administration and scoring of general cognitive screens belongs in the Psychometry, Assessment and Descriptive Psychopathology notes. Here the test is used only as part of differential reasoning. A score is a sample of performance under particular conditions; the diagnosis comes from the syndrome and trajectory in which that sample sits.

19.6 The cutoff trap: healthy controls are the wrong comparator

Separating depression from dementia is particularly difficult because both can produce genuine cognitive test abnormalities. The **Delayed Word Recall Test** illustrates the danger of importing a cutoff from the wrong comparison. A threshold of **fewer than 3 words out of 10** separated early dementia from healthy controls almost perfectly in the original comparison. When the same threshold was applied to dementia versus age-matched clinically depressed patients, **44%** of the depressed group were misclassified as having dementia.

The lesson is broader than this test. A threshold validated against healthy people answers a relatively easy question: who resembles a healthy control and who does not? Clinical practice asks a harder question: which of two disorders capable of impairing cognition better explains this patient? Performance that is unusual for health may be entirely non-specific when depression is the comparator. Diagnostic accuracy cannot be transported unchanged across populations merely because the test and cutoff remain the same.

TRAP **Do not equate "below a dementia-screening cutoff" with "therefore dementia."** A cutoff that performs impressively against healthy controls may perform poorly when the real alternative diagnosis is depression.

For an examination answer, state the comparator explicitly. For a clinical formulation, ask what the threshold was designed to discriminate and whether the patient resembles the population in which it was tested. The correct response to overlap is not to abandon cognitive testing. It is to stop asking a single score to perform the work of history, state examination, functional assessment and longitudinal review.

19.7 Longitudinal risk changes the meaning of "pseudodementia"

Patients thought to have depression-induced cognitive impairment later develop Alzheimer disease dementia almost **six times** as frequently as age-adjusted counterparts. This is a risk association, not a deterministic conversion rule. It does not mean that every depressed older adult with cognitive symptoms has latent dementia, and it does not justify naming a future diagnosis before the longitudinal evidence exists.

The figure matters because it changes follow-up from optional reassurance to part of diagnosis. If cognition improves with time and encouragement, that observation supports one formulation; if impairment persists into mood remission or a progressive pattern emerges, the formulation must be revised. The source evidence also associates depression with future vascular dementia, reinforcing that the relevant outcome is not restricted to a single dementia type.

Two opposite errors follow from ignoring this evidence. The first is false reassurance: declaring the cognitive syndrome "only depression" and ending cognitive surveillance after affective improvement. The second is premature closure in the other direction: treating a risk association as proof that dementia is already present. Longitudinal reasoning avoids both. It preserves therapeutic attention to depression while maintaining proportionate concern about persistent or progressive cognitive change.

19.8 Integrating the three Ds at the bedside and in an answer

Use a hierarchy, not a bag of signs. First determine whether the patient is in a stable state of consciousness and can sustain attention. Next reconstruct onset and course. Then examine the cognitive profile, testing context and response to encouragement. Finally, compare mood and cognition longitudinally. This order prevents an impressive memory deficit from distracting the examiner from a more urgent state disturbance and prevents a low score from being treated as a diagnosis.

DECISION QUESTION	FINDING THAT REDIRECTS REASONING	INTERPRETIVE CONSEQUENCE
Is consciousness altered?	Altered level with cognitive impairment	Prioritise delirium in the differential
Can attention be sustained?	Attention is prominently impaired with fluctuating alertness	Supports the delirium pattern
What is the temporal arc?	Acute or subacute and fluctuating versus insidious and slowly progressive	Separates state change from progressive decline
Does supportive testing matter?	Performance improves with enough time and encouragement	Supports depression-associated impairment, without proving it
What happens after mood recovery?	Cognitive impairment may persist, especially in memory and executive function	Continue longitudinal cognitive review
Was dementia already present?	New state change on a progressive baseline	Consider delirium superimposed on dementia

When findings point in different directions, weight them by discriminating value rather than counting them. Impaired episodic memory contributes little to the delirium-dementia separation because it may occur in both. By contrast, altered alertness, impaired attention and fluctuation form a coherent state-change pattern. Preserved orientation during an interview should also be weighted cautiously because orientation may remain intact early in dementia and a fluctuating delirium may be sampled during a clearer period. The question is which combination best explains the whole presentation, not which diagnosis collects the largest number of isolated signs.

Coexistence should be stated whenever it resolves an apparent contradiction. A progressive cognitive baseline does not cancel a new acute attentional disturbance; it makes delirium superimposed on dementia a necessary consideration. Similarly, a depressive episode does not make persistent cognitive impairment imaginary, and persistent impairment does not by itself reveal whether depression was prodromal, causally related or distinct. The formulation can legitimately contain more than one syndrome and more than one level of certainty.

A useful closing sentence identifies the next discriminating observation. For suspected delirium, that observation is the evolution of alertness and attention within the acute fluctuating course. For a dementia pattern, it is evidence of continued progressive decline from baseline. For depression-associated impairment, it is whether cognition improves under supportive conditions

and how it behaves as the mood episode remits. This makes follow-up part of diagnostic reasoning rather than an admission that the examiner failed to decide.

MNEMONIC

AAT: Arousal, Attention, Trajectory

Arousal asks whether consciousness is altered; **Attention** identifies the hallmark cognitive disturbance of delirium; **Trajectory** separates an acute fluctuating state, an insidious progressive decline and depression-associated cognitive impairment that requires follow-up rather than a promise of reversibility.

A concise written answer can therefore proceed in a disciplined sequence:

- Lead with consciousness, attention, onset and course, because these provide the strongest delirium-dementia separation.
- Add the shared and contrasting cognitive, perceptual, speech and behavioural features rather than presenting memory loss as unique to dementia.
- Close with coexistence and uncertainty: delirium can occur on dementia, depression-related impairment may persist, and a cutoff cannot replace longitudinal diagnosis.

Altered consciousness plus impaired attention and fluctuation points toward delirium; clear consciousness shifts the comparison toward dementia or depression, where trajectory, testing context and follow-up carry the distinction.

The final formulation should remain syndromic and revisable. Name the dominant pattern, state the evidence that supports it, acknowledge any baseline disorder, and identify what observation over time would change the conclusion. That is stronger than claiming certainty from memory impairment, a hallucination, an apparent response to encouragement or a single cutoff. The three Ds are differentiated most safely by combining state, domain and trajectory.

Assessment, biomarkers and the pharmacology of dementia

20 · Cognitive assessment in dementia: bedside batteries and neuropsychological profiles

Cognitive assessment should convert a vague complaint into a defensible pattern. The clinician must decide whether performance is genuinely abnormal, which domains are affected, whether the pattern fits the history and functional decline, and whether a brief screen is sufficient. A total score can support this reasoning, but cannot replace it. The marks lie in choosing the right level of assessment, interpreting performance in context, and describing a profile rather than announcing dementia from a single number.

Bedside testing and formal neuropsychological assessment answer related but different questions. A bedside battery offers a structured snapshot suitable for case finding, comparison over time

and communication between clinicians. Formal assessment examines the pattern in greater depth when the bedside result is equivocal, the presentation is unusual, or a domain-specific question matters. Item-level administration, standardised scoring and the generic psychometric properties of these instruments belong in the Psychometry, Assessment and Descriptive Psychopathology notes. Here the lens is dementia: selection, interpretation, serial use and syndrome-level profiling.

SCREEN	CONTEXT	PROFILE	COURSE
STRUCTURED SNAPSHOT	FAIR INTERPRETATION	DOMAIN PATTERN	SERIAL MEANING

20.1 What a bedside cognitive screen can and cannot establish

A screen is best treated as a structured sample of cognition. It can disclose a low score, an uneven pattern, a striking qualitative error or a discrepancy between the patient's confidence and performance. It can also provide a baseline for later comparison. It does not by itself establish cause, chronicity or functional consequence. Those conclusions require history from the patient and an informant, examination, assessment of everyday functioning and the wider diagnostic workup.

The **Mini-Mental State Examination** became the standard bedside dementia screen because it is regularly used, reproducible and clinically interpretable; its results correlate with histological change in Alzheimer disease. Among well-educated people whose scores are borderline, **10% to 25% may develop dementia within the next 2 years**. The point is predictive enrichment, not certainty. A borderline result identifies a group requiring clinical attention and follow-up; it does not identify which individual will progress.

RULE **Interpret the score as evidence within a formulation, never as the formulation itself.** State the instrument, testing context, domain pattern, functional history and longitudinal direction before drawing an etiological inference.

Before testing, clarify the question. Is the concern early decline despite apparently preserved conversation? Is the aim to document a baseline, follow an established disorder, explore an executive presentation, or decide whether formal testing is needed? The same brief instrument cannot be equally sensitive to every pattern. The question should determine the battery, rather than familiarity with a battery determining the question.

20.2 Interpreting the standard bedside screen in dementia

The total score has meaning only within the source's stated construct. The cited clinical summary gives a **maximum of 30 and describes scores of 20 or less as occurring with dementia, delirium, schizophrenia or affective disorder, alone or together**. Such low scores were not found in normal older people or in people with neurosis or personality disorder in that account. This is not a dementia-specific diagnostic cut-off. It is a broad correlation with markedly abnormal clinical populations. The commonly recalled alternative screening threshold and any India-specific adjustment are not supplied here and should not be substituted from memory.

Functional correlations help show why a total score must be translated back into life. In the same summary, **a total around 20 correlates with inability to keep appointments or use a**

telephone, while around 15 it marks inability to dress or groom. These associations illustrate increasing dependence; they do not permit the clinician to infer a person's exact abilities from the score. Function must still be asked about directly, because performance in a consultation and performance in a familiar environment are not interchangeable.

Serial trajectory may be more informative than an isolated result. If successive scores remain stable for **2 years**, the diagnosis of a relentlessly progressive Alzheimer process should be reconsidered. If they fall precipitously, rapidly progressive causes become more likely. In both situations, the correct response is diagnostic review, not mechanical restaging. Test conditions, intercurrent illness, sensory difficulty, language, mood and participation should be considered before attributing all change to neurodegeneration. Delirium itself is taught in the Stroke, TBI, delirium and focal brain syndromes notes.

GOOD TO REMEMBER

Record the qualitative behaviour beside the score: whether the patient encoded the material, benefited from cueing, lost set, produced intrusions, misunderstood instructions, became impulsive or showed disproportionate difficulty in one domain. That description often carries more diagnostic value than an unqualified total.

20.3 Ceiling effects, domain bias and false reassurance

A familiar screen may be "too easy" for the clinical question. The **ceiling effect** allows mild cognitive impairment and early dementia to escape detection. Age, education and language influence the score. Heavy dependence on orientation, memory and language means that executive dysfunction may be under-sampled, including presentations in which frontal systems are central. The same language dependence makes the instrument less adequate when subcortical white-matter damage is the principal problem.

Domain dependence is especially important when the suspected profile is frontal-executive rather than amnestic. Vascular cognitive impairment literature illustrates the problem: a patient may struggle with planning, set shifting and cognitive control while retaining enough orientation and language to appear reassuring on a language-heavy total. The remedy is not to reinterpret the same total more aggressively. It is to supplement the examination with an instrument that samples the domain of concern and to examine the behaviour displayed during testing.

In Parkinson disease cognition, the **Montreal Cognitive Assessment** is favoured over the standard screen because it covers more cognitive subdomains, has less ceiling effect and is more sensitive to longitudinal change. This is a context-specific illustration of battery selection, not permission to teach Parkinson neuropsychiatry here. That broader subject belongs in the Neuropsychiatry of systemic and neurological disease notes.

TRAP **A high bedside score does not exclude meaningful cognitive disorder.** The classic error is to quote the total, ignore the presenting domain and stop. Early disease, executive syndromes, high premorbid ability and a ceiling-prone instrument can all make the number falsely reassuring.

20.4 Choosing among bedside batteries

The **broader multidomain screen** is more sensitive than the standard examination for mild impairment. It is particularly useful when memory is the only reported symptom or when concern persists despite a high standard-screen score. It is widely used across neurodegenerative dementias and detects early cognitive impairment in Parkinson disease better than the standard screen. Its stronger sampling of frontal dysfunction also makes it more suitable than a brief mental-status instrument when mild cognitive impairment is suspected and the clinical question is not purely orientation and recall.

Choice should balance sensitivity, domain coverage, testing burden, language and educational fit, sensory or motor limitations, examiner competence, permission to use the instrument, and the purpose of serial follow-up. The batteries are not interchangeable labels for "cognition". Their architecture determines what they are likely to reveal.

BEDSIDE OPTION	DOMAIN EMPHASIS IN THE CITED COMPARISON	BEST DEMENTIA-ASSESSMENT USE	MAIN CAUTION
Standard mental-status screen	Orientation, verbal registration and recall, attention, language, praxis and visuospatial performance	Familiar broad snapshot and serial comparison	Ceiling and executive-domain limitations
Broader multidomain screen	Orientation, attention, memory, naming, fluency, repetition and visuospatial-executive performance	Mild or executive concern despite a reassuring standard score	Education-related false-positive classification
Addenbrooke cognitive battery	Attention and orientation, memory, fluency, language and visuospatial performance	Early Alzheimer pattern, Alzheimer versus frontotemporal pattern, separation of organic brain disease from psychiatric states, and selected parkinsonian cognitive syndromes	Use only validated administration and scoring material
Cambridge cognitive battery	Orientation, language, memory, attention, praxis, calculation, abstraction and perception	A more extended generic screen when broader sampling is needed	Greater time burden than briefer bedside options

The standard and broader screens occupy a briefer bedside slot, while the revised Addenbrooke and Cambridge batteries are more extended. Administration time should therefore be selected deliberately rather than treated as an inconvenience discovered midway through testing. A hurried, poorly heard or linguistically mismatched administration is not made valid merely because it finishes quickly.

The **shorter validated Addenbrooke form** can be completed in **under 5 minutes and is scored out of 30**. Publicly available alternatives include the Mini-Cog, St. Louis Mental Status

Examination, Kokmen Short Test of Mental Status and Rowland Universal Dementia Assessment Scale. Their importance is practical because the Mini-Mental State Examination and Montreal Cognitive Assessment are no longer in the public domain and have been monetised. Availability, licensing and authorised materials therefore belong in test selection. Convenience does not justify copying protected item content, and this section deliberately describes rather than reproduces it.

20.5 Language, education and the Indian consultation

Every score is produced by a patient, an examiner, a language and a setting. Education can influence familiarity with test-taking, abstraction, paper-and-pencil demands and the strategies used to solve unfamiliar tasks. Language affects comprehension, lexical access and the meaning of apparently simple instructions. Sensory impairment, fatigue, anxiety, motor slowing and poor rapport can further lower performance without having the same diagnostic significance as neurodegeneration.

The broader screen samples more domains, but **low-education misclassification** is a recognised risk: cognitively normal people with limited educational attainment may be labelled as having dementia. The solution is not to waive concern whenever schooling is limited. It is to interpret performance alongside premorbid opportunity, occupational and daily competence, informant evidence, observed decline and the suitability of the chosen instrument. A decline from the person's own prior level is clinically different from lifelong unfamiliarity with the task.

IN INDIA

Do not import an education-insensitive threshold into a multilingual, variably schooled consultation and call the result objective. Test in the most appropriate available language, document the language and educational context, check hearing and vision, and prefer a validated culturally suitable instrument where available. When fit is doubtful, describe the observed domain pattern and escalate the assessment rather than manufacturing precision.

Translation at the bedside is not neutral. Changing a word, adding an explanation or allowing a relative to coach the patient may alter task difficulty. If an adaptation is necessary, document it and avoid pretending that the resulting total is strictly equivalent to a standard administration. A qualitative account can remain clinically useful even when a normative inference is unsafe.

MNEMONIC

SNAP

Select the battery for the question; **N**ote language, education, sensory and testing context; **A**ssess the pattern across relevant domains; **P**lace the result beside function and progression. The mnemonic prevents a bedside score from becoming a diagnosis by reflex.

20.6 When formal neuropsychological assessment adds value

Formal assessment is most useful when it answers a defined question. Referral is justified when bedside findings and history disagree, decline is subtle but consequential, the pattern is atypical,

executive or language dysfunction requires deeper characterisation, premorbid ability complicates interpretation, or serial measurement must be more domain-sensitive. It can also help examine whether depression or non-credible performance contributes to apparent impairment without reducing that judgement to a single performance sign. Depressive pseudodementia and the broader depression-dementia differential require their own clinical assessment and are not reproduced here.

A useful referral specifies the suspected domains, onset and course, functional change, language, education, occupation, sensory and motor constraints, psychiatric symptoms, medication context and the decision that the result will inform. “Rule out dementia” is too blunt. The neuropsychologist can better select and interpret measures when asked whether the profile is amnestic, dysexecutive, language-led, visuospatial, patchy, or globally slowed, and whether observed weaknesses are consistent with the history.

- **Define the decision:** diagnosis, differential pattern, baseline, progression or capacity for a particular everyday task.
- **Describe the trajectory:** onset, tempo, fluctuation and the relation between complaint and observed decline.
- **State the context:** education, language, occupation, sensory or motor limitation and relevant psychiatric state.
- **Request a profile:** strengths, weaknesses, error types, cueing response and coherence across domains.
- **Reintegrate the result:** compare it with function, examination, investigations and longitudinal observation.

Formal testing remains an examination under conditions, not a direct image of pathology. A profile may support an etiological hypothesis, reveal internal inconsistency or show that a bedside total concealed an important deficit. It should then be reintegrated with the clinical syndrome. Neuropsychological batteries taught as instruments, including their administration and scoring, remain with the Psychometry, Assessment and Descriptive Psychopathology notes.

20.7 Neuropsychological profiles across common dementia syndromes

The profile is a pattern of relative strengths, weaknesses and sequence. Etiological subtyping describes Alzheimer disease as typically memory-led, while frontotemporal lobar degeneration is typically led by speech, behaviour or executive dysfunction rather than by the same early amnestic emphasis. These are characteristic tendencies. They organise the differential but do not prove pathology in an individual patient.

CLINICAL PATTERN	LEADING COGNITIVE OR BEHAVIOURAL FEATURES	INTERPRETIVE CLUE
Alzheimer pattern	Early short-term memory and language impairment, followed by executive and then visuospatial dysfunction	Progressive amnestic-language beginning with later spread across domains
LATE pattern	An amnestic syndrome resembling the Alzheimer presentation	Neuropsychology may show resemblance rather than settle the underlying pathology
Frontotemporal pattern	Relative memory sparing, with language impairment or executive dysfunction predominating	Look beyond recall totals to behaviour, language and cognitive control
Vascular dementia pattern	Patchy deficits shaped by the location of vascular lesions	Uneven performance is more informative than forcing a single smooth sequence
Lewy body pattern	Visuospatial, attention and executive deficits more prominent than memory impairment, with cognitive fluctuation	Visual hallucinations, spontaneous parkinsonism, earlier psychosis or personality change, and REM-sleep disturbance strengthen the clinical pattern
Mild cognitive impairment pattern	Reduced processing speed and choice-reaction performance, with mild inconsistent forgetfulness	Cognitive symptoms occur without the functional impairment that defines dementia in this comparison

The behavioural variant of frontotemporal dementia sharpens the contrast. Patients may lack awareness of inappropriate behaviour, show relatively preserved memory with benefit from cues, and retain general visuospatial ability despite later cognitive deficits in contrast to the rapid memory loss described in the Alzheimer comparison. Cueing response is therefore not a decorative detail: it helps distinguish failure to retrieve efficiently from a more severe failure to retain information.

The table is a set of complementary clinical patterns, not a canonical map covering every domain for every dementia. Mixed pathology, disease stage, premorbid ability and focal presentations can blur the expected shape. Use the profile to generate and test hypotheses, then reconcile it with the course, neurological findings, function and biomarker or imaging workup described in their respective sections.

20.8 Cortical versus subcortical pattern and final synthesis

The **cortical-subcortical distinction** remains a useful descriptive scaffold. A cortical pattern shows module-specific deficits: amnesia, aphasic features and apraxia, while alertness and executive function may be preserved early. A subcortical pattern is characterised by slowing, impaired attention and executive function, reduced verbal output, and forgetfulness that improves with cueing; language structure and praxis remain relatively preserved. The distinction is about the organisation of deficits, not merely their severity.

This contrast explains why a language-heavy bedside screen can miss clinically important subcortical dysfunction. A patient may be slow, inattentive and dysexecutive yet retain naming, praxis and basic language structure. Conversely, a cortical syndrome may show a striking focal failure even when general alertness appears normal. Huntington disease is a standard subcortical example, but its disease-specific neuropsychiatry belongs in the Neuropsychiatry of systemic and neurological disease notes.

PITFALL

Do not equate poor recall with one mechanism. Ask whether material was encoded, retained, retrieved spontaneously or recovered with cueing. Amnesia, inefficient retrieval, inattention and language failure can all produce "memory difficulty", but they imply different profiles.

The final synthesis should read as a compact argument. Name the battery and why it suited the question. Describe testing conditions and any limitation. Report the total only within its proper construct. State the affected and relatively preserved domains, qualitative errors and cueing response. Link the pattern to everyday function and the longitudinal course. Then say which etiological pattern it supports, what it fails to explain, and whether formal assessment or another investigation is needed.

This is also the safest examination strategy. A candidate who lists instruments demonstrates recall; a candidate who explains selection, context, profiles and limitations demonstrates clinical reasoning. Bedside testing begins the cognitive formulation. Neuropsychological profiling refines it. Neither replaces the history of decline or the evidence of functional consequence.

A cognitive score screens; the domain pattern, functional decline, testing context and serial course make it clinically meaningful.

21 · Structural and functional neuroimaging in dementia (MRI, FDG-PET, amyloid-PET, DAT scan)

Neuroimaging earns marks when it answers a clinical question, not when it becomes a catalogue of scans. In a patient with cognitive decline, imaging is used to identify the anatomical pattern, estimate vascular or degenerative burden, look for a functional network signature, or test for a molecular or dopaminergic signal. The report becomes meaningful only after it is reconciled with the history, examination and cognitive phenotype. A scan can strengthen or weaken a diagnostic formulation; it should not replace one.

The useful framework separates structural MRI, FDG-PET, amyloid-PET and tau-PET, and dopamine-transporter imaging. Structural imaging asks where tissue loss or vascular injury is distributed. Metabolic or perfusion imaging asks which networks are functioning abnormally. Molecular imaging asks whether a targeted pathological signal is present. Dopamine-transporter imaging contributes to a narrower differential involving disorders in which that signal is abnormal. The physics, tracer kinetics and acquisition principles belong in the Neurochemistry,

Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes; the concern here is dementia-specific interpretation.

MRI	FDG	AMYLOID	DAT
ANATOMICAL PATTERN	NETWORK FUNCTION	PATHOLOGICAL SIGNAL	DOPAMINERGIC SUPPORT

21.1 Start with the question the scan must answer

A useful request begins with the unresolved differential. “Cognitive decline” is not enough. The clinician should state whether the dominant uncertainty concerns a degenerative pattern, vascular contribution, frontotemporal distribution, an occipital functional pattern, or a dopaminergic syndrome. This improves the chance that the radiological description is translated into a clinically usable answer. It also prevents the common habit of ordering progressively more specialised tests merely because the preceding result was non-specific.

Before ordering, write down the expected consequence of each possible result. If an amyloid-negative study would meaningfully weaken a proposed AD aetiology, the test has a clear role. If a characteristic functional pattern would resolve a genuine degenerative differential, the request is similarly focused. If neither a positive nor a negative result would alter formulation, counselling or further assessment, the request is poorly framed. This pre-test discipline matters because imaging abnormalities are easy to notice but harder to interpret responsibly. The more specialised the test, the more precisely the clinician should define the probability shift being sought.

MODALITY	PRIMARY CLINICAL QUESTION	BEST USE IN REASONING	MAIN INTERPRETIVE CAUTION
MRI	Where is atrophy or vascular injury distributed?	Match anatomy to the cognitive and behavioural phenotype	Atrophy can support a formulation without being diagnostic by itself
FDG-PET or perfusion imaging	Which cortical or subcortical networks show reduced function?	Compare the observed pattern with characteristic dementia patterns	A pattern must remain concordant with structural imaging and phenotype
Amyloid-PET or tau-PET	Is the targeted molecular signal present, and does it fit the suspected process?	Use molecular evidence to shift the probability of an aetiological diagnosis	Do not turn one biomarker into a complete clinical diagnosis
DAT scan	Is dopamine-transporter imaging abnormal in a relevant syndrome?	Support a dopaminergic differential when clinical uncertainty remains	An abnormal result is shared by several neurological disorders

RULE **Ask anatomy, function, pathology or dopamine-transporter status as separate questions.** Choosing the modality follows from the question; interpreting it requires convergence with the phenotype.

This sequence is also a safeguard against overtesting. Structural and functional findings often answer different parts of the same problem. A mismatch should prompt re-examination of the clinical formulation, image quality, disease stage and competing pathology. It should not be "resolved" by automatically privileging the more technologically impressive scan. In examination answers, state the diagnostic question first, describe the expected distribution next, and finish with the limitation.

The radiology request and the final formulation should use the same vocabulary. If the request asks about a frontotemporal process, the answer should address frontal and anterior-temporal distribution rather than merely stating that atrophy is present. If vascular contribution is the issue, burden, territory and relationship to the clinical deficits matter more than an undifferentiated label of small-vessel change. A precise question invites a pattern-based report; a vague question invites a list of incidental observations.

21.2 Structural MRI: pattern, burden and concordance

MRI contributes by showing the distribution of atrophy, ventricular enlargement and white-matter abnormality. The central interpretive distinction is between a supportive pattern and a stand-alone diagnosis. In AD, MRI findings are **informative rather than diagnostic**. The same principle applies more broadly: the image should fit the syndrome's dominant cognitive, behavioural or neurological features before it is allowed to carry substantial diagnostic weight.

In the amnestic degenerative pattern, medial-temporal atrophy is the structural emphasis. Even at a mild clinical stage, the cited quantitative description reports **20-30%** loss of entorhinal cortex volume and **15-25%** loss of hippocampal volume, together with ventricular enlargement. These figures illustrate an important clinical point: apparently mild symptoms need not imply mild underlying tissue injury. They should not, however, be converted into a bedside cut-off or used as a substitute for formal image interpretation.

Vascular disease requires attention to burden and distribution rather than a reflex equation of any white-matter change with vascular dementia. The supplied description associates vascular dementia with progressively increasing **white-matter hyperintensity burden**. The clinician must still ask whether the imaging pattern plausibly explains the cognitive profile and course. Patchy injury may produce uneven deficits, while diffuse subcortical network involvement may produce a different clinical emphasis. Imaging-clinical concordance is therefore more valuable than the mere presence of an abnormality.

TRAP **Do not reproduce a summary-table error as though it were settled anatomy.** One cited table gives a **summary-table medial-temporal entry for FTD**, yet the same source's detailed disease account describes **frontal and anterior-temporal atrophy**, particularly involving prefrontal cortex, insula, anterior cingulate, striatum and thalamus, with greater involvement of the non-dominant hemisphere. State the inconsistency, then use the detailed syndrome-specific pattern for clinical reasoning.

The conflict is exam-relevant because it tests whether the candidate understands a pattern or merely recalls a row. A behavioural frontotemporal presentation should be examined for a frontal

and anterior-temporal structural distribution, not forced into an undifferentiated medial-temporal template. Conversely, medial-temporal change in a patient with a different phenotype should not erase the clinical syndrome. Mixed or overlapping pathology remains a clinical possibility, but it cannot be declared from an isolated image description.

Read the structural study in a fixed order. Decide whether volume loss is focal or diffuse, symmetrical or asymmetrical, and predominantly medial-temporal, frontal or anterior-temporal. Then assess whether vascular abnormalities are patchy, confluent or anatomically related to the deficits under discussion. Finally, compare these observations with ventricular size and the overall degree of tissue loss. This order converts descriptive radiology into clinical neuroanatomy. It also exposes discordance early: a striking behavioural syndrome with an unrelated structural distribution should trigger reconsideration rather than a forced match.

PITFALL

Do not diagnose vascular dementia from incidental white-matter change alone. First ask whether the amount, location and network effect of the vascular injury are sufficient to account for the observed decline, and whether the clinical course is concordant.

21.3 Functional imaging: recognise distributions, not bright colours

Functional imaging is most useful when it converts a vague impression of “reduced uptake” into a distributed network pattern. FDG-PET is interpreted as a regional metabolic pattern, while perfusion imaging offers a related pattern-based differential. The acquisition physics is not the task here. The task is to identify the dominant regions, decide whether the distribution is coherent, and compare it with the structural and clinical picture.

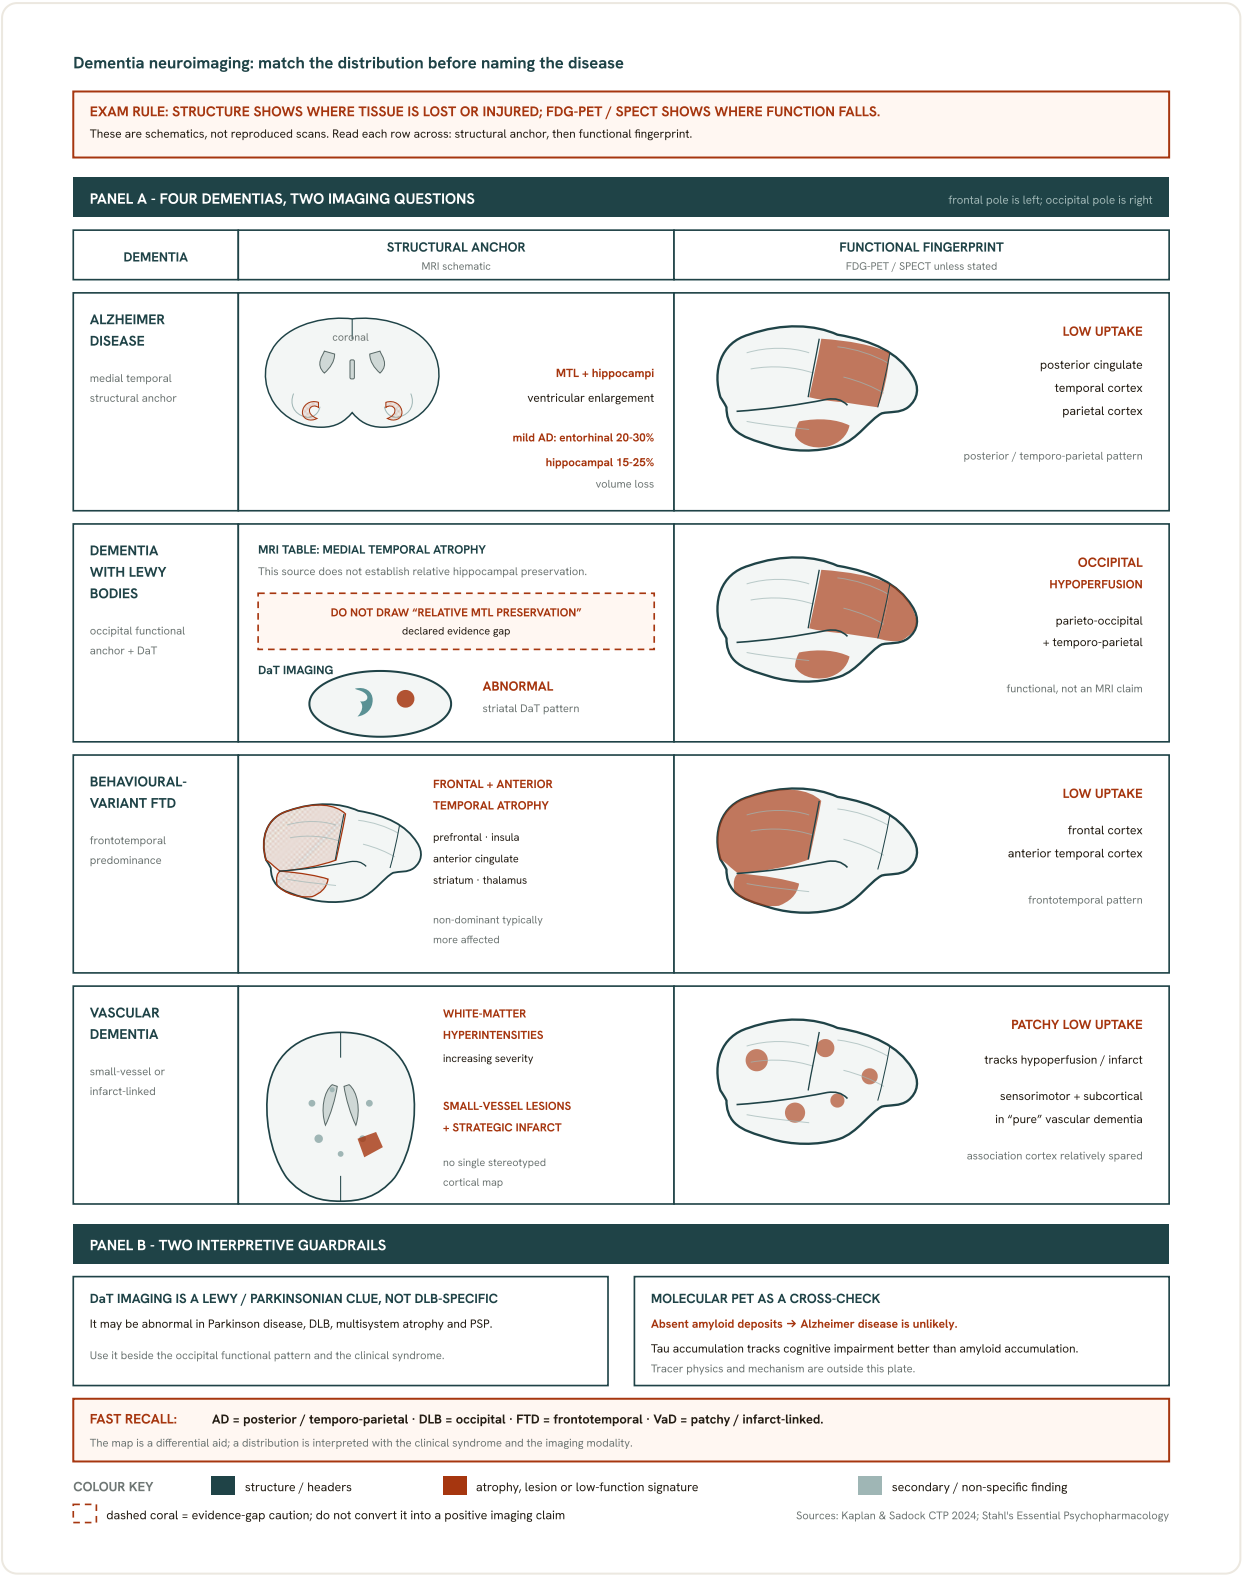


Fig 32.8 Neuroimaging patterns in the dementia differential. Alzheimer disease centres on medial temporal and hippocampal atrophy with posterior cingulate, temporal and parietal functional reduction; DLB adds an occipital functional pattern and abnormal DaT imaging; behavioural-variant F'TD is frontal and anterior temporal; vascular dementia combines white-matter and infarct burden with patchy, lesion-linked functional loss. The DLB schematic does not claim relative hippocampal preservation because that comparison was not grounded in the supplied evidence.

The accompanying figure should be read as a schematic of relative patterns, not as a substitute for viewing or formally reporting patient images. The following table condenses the high-yield differential.

CLINICAL SYNDROME	CHARACTERISTIC FUNCTIONAL PATTERN	HOW TO USE IT
AD	**Posterior cingulate, temporal and parietal reduction**	Supports the characteristic posterior-temporoparietal network pattern when concordant with the phenotype
Vascular dementia	**Patchy reduction corresponding to hypoperfusion or infarct location**	Relate each deficit to the vascular distribution rather than expecting a smooth degenerative pattern
FTD	**Frontal and anterior-temporal reduction**	Seek concordance with behavioural or language-led frontotemporal dysfunction
DLB	**May show occipital hypoperfusion or reduced function**	Use as supportive functional evidence within the full clinical differential

A second description of relatively "pure" vascular dementia adds a **sensorimotor and subcortical pattern with relative association-cortex sparing**, contrasting with the posterior or temporoparietal emphasis of AD. This is not a contradiction with a patchy vascular pattern. It reflects the fact that vascular cognitive syndromes can be expressed through the territory and network disrupted. In an answer, link the functional abnormality to the vascular anatomy instead of presenting vascular dementia as a single uniform metabolic map.

MNEMONIC

Back, Broken, Behaviour, Back-of-head

For the functional differential: AD is predominantly **back** in posterior cingulate and temporoparietal regions; vascular dementia is **broken** or patchy by lesion distribution; FTD follows **behaviour** and language networks in frontal and anterior-temporal regions; DLB may point to the **back of the head** through an occipital pattern.

Pattern reading should proceed in layers. First identify whether the reduction is focal, patchy or network-like. Next name the dominant lobes and whether association cortex appears relatively spared. Then compare the functional map with MRI. Finally, ask whether the patient's phenotype is expected from that network. This disciplined sequence prevents a colourful image from acquiring more authority than its clinical context permits.

A single affected region is usually less informative than the relationship among regions. Posterior cingulate involvement gains meaning when temporal and parietal reduction forms a coherent posterior network. Frontal reduction gains meaning when anterior-temporal involvement and the clinical syndrome point in the same direction. Occipital change gains meaning within the DLB differential rather than as an isolated visual label. Vascular reduction gains meaning through correspondence with hypoperfusion or infarct distribution. In each case, the pattern is the teaching unit. Memorising one lobe without its network context invites false certainty.

21.4 Amyloid-PET and tau-PET: molecular evidence changes probability

Molecular PET addresses pathology more directly than a regional atrophy or metabolism pattern, but its interpretation remains probabilistic. The most important negative inference is that **absence of amyloid deposits** indicates that dementia is not likely to be due to AD. Notice the direction of the statement: a negative amyloid study argues against that aetiology. The converse should not be carelessly rewritten as "a positive scan proves the entire clinical diagnosis." Clinical syndrome, functional effect and alternative explanations still matter.

Tau-specific PET detects hyperphosphorylated tau aggregates described across AD, progressive supranuclear palsy, corticobasal degeneration, frontotemporal lobar degeneration and chronic traumatic encephalopathy. For the candidate, the discriminating point is not merely that amyloid and tau can each be imaged. It is that the degree of tau accumulation correlates better with cognitive impairment than the degree of amyloid accumulation. This **tau burden-cognition relationship** helps explain why a pathological signal and current clinical severity are related but not interchangeable constructs.

The proper mental model is a chain rather than a verdict. Molecular signal addresses the presence of targeted pathology; structural imaging addresses accumulated anatomical effect; functional imaging addresses network dysfunction; clinical assessment addresses lived cognitive and functional consequence. Agreement across these layers increases coherence. Discordance is information: it may indicate timing, co-pathology, a poorly matched phenotype, or a limitation of the test. Without a supported threshold or performance figure, do not invent one in an examination answer.

Negative results also have unequal meanings. The supplied evidence gives a clear negative inference for amyloid imaging, but it does not provide equivalent exclusion rules for every other modality. A non-characteristic FDG pattern, for example, should be described as non-supportive or discordant rather than automatically treated as proof of an alternative dementia. The disciplined answer uses the strongest direction actually established for each test and leaves unsupported reverse inferences alone. This is especially important when a candidate remembers a familiar slogan but cannot state the population, comparator or threshold behind it.

> **GOOD TO REMEMBER**
>
> Read a negative amyloid study in the direction actually supported: it makes AD an unlikely cause of the dementia. Do not reverse that statement into an unqualified claim that any positive study establishes symptomatic AD.

21.5 DAT scan: a supportive test with a deliberately narrow claim

Dopamine-transporter imaging is abnormal in **Parkinson disease, DLB, multisystem atrophy and progressive supranuclear palsy**. Therefore, an abnormal scan is not unique to DLB and cannot by itself separate all disorders in that group. Its value is narrower: in the appropriate clinical differential, it supports the presence of an abnormal dopaminergic imaging signal and should be integrated with the dementia phenotype and neurological examination.

This is a useful example of why “abnormal” and “diagnostic” are not synonyms. A test can discriminate one broad mechanism or disease family while remaining unable to identify a single member of that family. The result must be interpreted against the question asked. If the clinical problem is not a differential in which dopamine-transporter status would change probability, the scan adds little merely by producing another abnormality.

Do not import the neuropsychiatry or treatment algorithms of Parkinson disease into this discussion. Those belong in the Neuropsychiatry of systemic and neurological disease notes. Here, retain only the dementia-imaging point: DLB may show an occipital functional pattern, and dopamine-transporter imaging can be abnormal, but neither statement is a substitute for the full syndrome diagnosis.

21.6 Integrating discordant results

Discordance is a prompt to think, not a reason to average the reports. A structural pattern may be subtle while functional change is convincing; molecular signal may be present without a matching clinical phenotype; vascular injury may coexist with a degenerative distribution. The clinician should identify exactly which layer disagrees and then reconsider the working hypothesis. “Mixed dementia” should not become a wastebasket label for every mismatch.

A practical integration sequence is:

- Restate the dominant clinical syndrome and the unresolved aetiological question.
- Describe MRI by distribution and burden, then ask whether the anatomy explains the phenotype.
- Describe functional imaging by network pattern, not merely as reduced uptake.
- Use molecular or dopamine-transporter imaging only for the probability shift that the test can legitimately provide.

When reports disagree, return to the images and the referral question where possible. Clarify whether the report is describing a definite pattern, a mild non-specific change, or a technically limited study. Compare prior imaging when available to understand progression, but avoid claiming a rate or interval that is not established. The clinical conclusion should state what the imaging supports, what it argues against, and what remains unresolved.

Fluid biomarkers can provide another biological layer, but their dementia-specific directions and emerging blood assays are taught in the dedicated biomarker section. Likewise, cognitive batteries contribute phenotype and severity, while their administration and general scoring belong in the Psychometry, Assessment and Descriptive Psychopathology notes. Imaging is strongest when it is one coherent part of this multimodal formulation.

21.7 Writing the examination answer

A high-scoring answer is organised by modality and clinical purpose. Begin with structural MRI, emphasising distribution of atrophy, ventricular enlargement and vascular white-matter burden. State that MRI is supportive rather than independently diagnostic. Then give the functional differential: posterior cingulate-temporoparietal reduction in AD, patchy lesion-linked reduction

in vascular dementia, frontal and anterior-temporal reduction in FTD, and a possible occipital pattern in DLB. Add the molecular PET principle that absence of amyloid makes AD unlikely and that tau burden tracks cognitive impairment more closely than amyloid burden. Finish with the narrow DAT point and its lack of uniqueness.

The best answer also contains its limitations. It does not teach PET or MRI physics, invent cut-offs, equate every white-matter lesion with vascular dementia, or conceal the conflict in the cited FTD MRI table. It distinguishes a characteristic pattern from a diagnostic criterion and a supportive biomarker from a complete diagnosis. That combination of recall and restraint is the real assessment target.

> Memorise the distributions: AD is posterior cingulate-temporoparietal, vascular dementia is patchy or lesion-linked, FTD is frontal and anterior-temporal, and DLB may be occipital; every pattern remains supportive and phenotype-dependent.

22 · CSF and emerging blood biomarkers (amyloid-beta-42, total and phospho-tau, plasma p-tau217)

Fluid biomarkers answer a biological question that the clinical interview cannot answer alone. They indicate whether a patient's cognitive syndrome is accompanied by a pattern of amyloid deposition, tau pathology or neuronal injury. Their value is greatest when the phenotype, structural assessment and longitudinal course leave a genuine question about the underlying disease process. They do not replace history, informant evidence, examination, cognitive assessment or the search for alternative causes.

The core examination task is to recognise the direction and meaning of the principal cerebrospinal-fluid changes, then place them within the **AT(N) research framework**. The same logic explains why blood testing is attractive: a less invasive sample could make biological assessment easier to deploy. However, an emerging marker is not automatically a stand-alone diagnosis, and a named analyte must not be given an invented threshold or accuracy estimate.

22.1 What a dementia fluid biomarker contributes

A biomarker is evidence about biology, not a synonym for dementia. The clinician begins by establishing whether there is an acquired cognitive decline, determines its pattern and severity, assesses functional consequences, and looks for competing neurological, psychiatric, medical and drug-related explanations. Fluid results then refine the causal formulation. This sequence prevents a laboratory pattern from being mistaken for the clinical syndrome itself.

The useful mental model is to separate the biological questions. Is there evidence of amyloid pathology? Is there evidence of pathological tau aggregation? Is there evidence of **neurodegeneration or neuronal injury**? These questions are related but not interchangeable. A result may speak strongly to an axis while saying little about another. Therefore, the report should be read analyte by analyte before it is compressed into an overall impression.

- **Phenotype:** What cognitive and behavioural syndrome is present, and how convincing is the evidence of decline?
- **Biology:** Which pathological axis does each measured analyte represent?
- **Concordance:** Do the clinical formulation and the biomarker pattern support the same explanation?

RULE **Diagnose the patient, then interpret the biomarker.** A fluid result modifies the probability of a biological process within a clinical formulation; it does not independently establish the patient's syndrome, severity or care needs.

This distinction also clarifies consent. Before testing, the clinician should know what uncertainty the sample is intended to resolve and how a positive, negative or discordant result would alter counselling or further assessment. Testing without a decision question creates information without a plan. It can also produce misplaced certainty when the report is read outside the context in which it was requested.

22.2 The core cerebrospinal-fluid profile

In **Alzheimer disease (AD)**, the characteristic amyloid signal is **low CSF amyloid-beta-42** relative to healthy individuals. The corresponding **CSF amyloid-beta-42/amyloid-beta-40 ratio** is also low. The ratio expresses the disease-relevant amyloid analyte relative to the broader amyloid pool rather than treating the absolute concentration in isolation. For examination purposes, both belong to the amyloid axis and move downward.

The tau side moves in the opposite direction: **CSF total tau** and **CSF phosphorylated tau** are increased relative to healthy individuals. Their shared upward direction should not erase their conceptual difference. Phosphorylated tau represents the tau-pathology axis in the biological framework, whereas total tau is placed with neurodegeneration or neuronal injury. This is why collapsing both measures into a generic tau label earns less than identifying which measure is being discussed and what biological question it answers.

FLUID MEASURE	DIRECTION IN AD	FRAMEWORK MEANING	INTERPRETIVE ROLE
Absolute amyloid measure	Decreased	Amyloid pathology	Read as an amyloid signal, not as a measure of clinical severity
Amyloid ratio	Decreased	Amyloid pathology	Frames the disease-relevant analyte relative to the accompanying amyloid measure
Phosphorylated tau	Increased	Aggregated tau	Answers the tau-pathology question
Total tau	Increased	Neurodegeneration or neuronal injury	Supports an injury signal but is not, by itself, disease-specific

The directions are more examinable than any remembered concentration. A numerical result only acquires meaning through the laboratory's validated method, units, quality controls and reference interpretation. The allowed evidence does not provide universal fluid cut-offs, so none should be memorised from this fragment. The safe answer is the pattern: amyloid measures are low, while both tau measures are high, with phosphorylated and total tau assigned to different biological axes.

Absolute concentration and biological classification also belong at different levels of reasoning. The analyser produces a concentration or ratio; the laboratory compares that result with its validated interpretive system; the clinician decides whether the resulting axis assignment is coherent with the patient. Skipping any of these levels invites error. A raw value copied into a case summary without units or laboratory interpretation is incomplete, while a categorical label copied without the underlying analytes conceals how the conclusion was reached.

MNEMONIC

Amyloid sinks; tau rises

For the core CSF pattern, the amyloid concentration moves down, the amyloid ratio moves down, and total tau and phosphorylated tau move up. Then separate phosphorylated tau as the tau-pathology signal from total tau as the neuronal-injury signal.

The must-memorise CSF pattern is low amyloid concentration and low amyloid ratio, with high total tau and high phosphorylated tau.

22.3 Reading the biological framework correctly

The framework sorts biomarkers into **three** biological axes. It was proposed by the **NIA-AA Research Framework** in **2018**. The conventional abbreviations are **A** for aggregated amyloid, **T** for aggregated tau and **(N)** for neurodegeneration or neuronal injury. The parentheses around the injury axis are meaningful in the notation: injury evidence is useful for characterising the biological state, but it is not specific to the defining protein pathologies.

A	T	(N)
AMYLOID	TAU PATHOLOGY	NEURONAL INJURY

For the fluid panel, the absolute amyloid measure or its ratio supplies A; phosphorylated tau supplies T; and total tau supplies (N). The framework also accommodates imaging correlates on these axes, but imaging selection and interpretation belong with the structural and functional neuroimaging discussion. The physics of those modalities is covered in the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

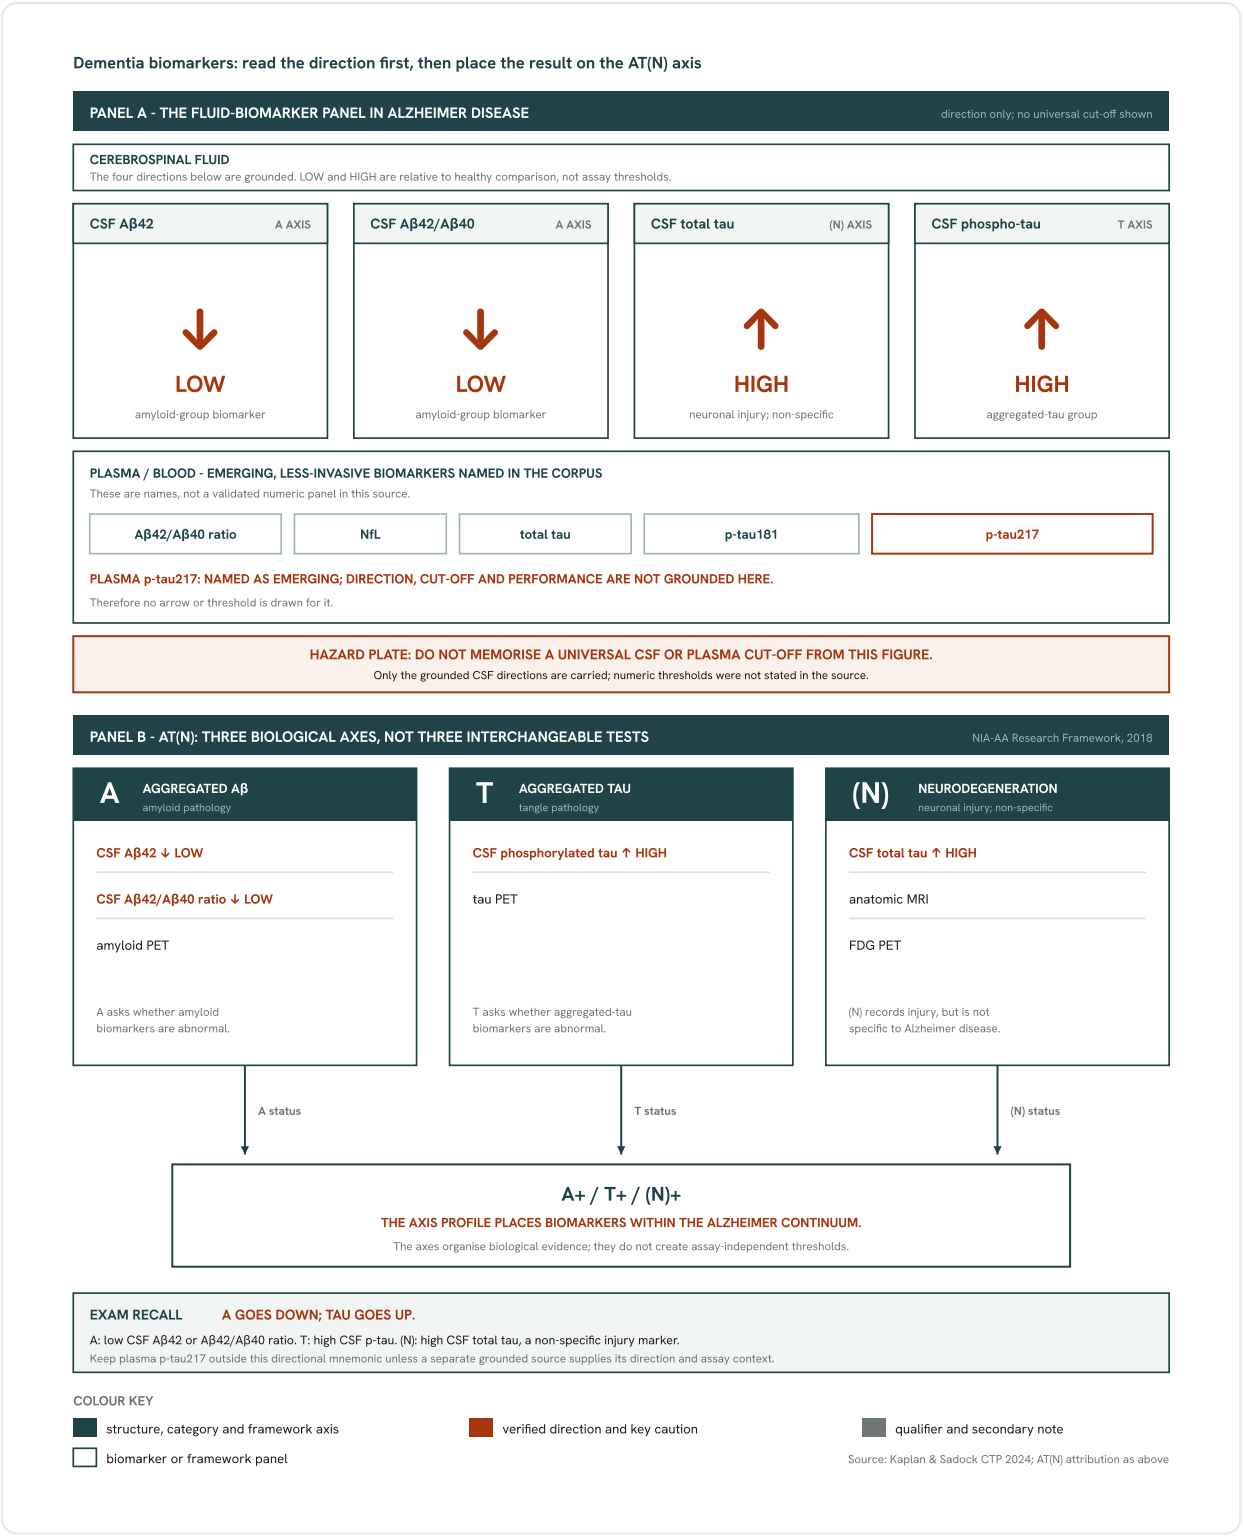


Fig 32.9 Dementia fluid biomarkers and the AT(N) framework. In Alzheimer disease, CSF Aβ42 and the Aβ42/Aβ40 ratio are low, while CSF total tau and phosphorylated tau are high. Plasma p-tau217 is shown only as an emerging named biomarker because the grounded source supplies no direction, cut-off or performance estimate.

The accompanying figure should be used as a translation device. Begin with the analyte, ask which axis owns it, and only then state the direction or implication. This prevents the common mistake of treating all abnormal biomarkers as equivalent evidence. A positive injury signal does not automatically identify the protein pathology producing that injury. Conversely, an amyloid-axis result and a tau-axis result answer more biologically specific but still distinct questions.

The notation can describe combinations of positive and negative axes, but classification should not become a substitute for clinical reasoning. A shorthand combination is a compact biological summary. It does not convey the pattern of memory loss, language impairment, behavioural change, disability, comorbidity or the wishes of the patient and family. Those dimensions still determine whether the finding is clinically coherent and what should happen next.

The framework is best treated as a coordinate system, not a severity ladder. Each axis tells the reader which kind of biology has been detected; it does not, merely by its place in the notation, rank functional impairment. This protects against a subtle but common overreach: turning a research classification into an unsupported statement about prognosis, rate of decline or current decision-making capacity. Those judgments require their own clinical evidence and should remain visible in the formulation.

TRAP **Putting both tau measures under T.** Phosphorylated tau belongs to the aggregated-tau axis, whereas total tau is assigned to (N), the neurodegeneration or neuronal-injury axis. Their concentrations move in the same direction in AD, but their framework meanings differ.

22.4 Pattern interpretation and discordance

A fluid panel should be interpreted as a pattern rather than as a row of isolated flags. Start with the amyloid axis, then the tau-pathology axis, then the injury axis. After that, return to the clinical question. A concordant pattern may strengthen a biological formulation; a discordant pattern should trigger review rather than creative arithmetic.

Discordance has several possible levels. The analytes may disagree with each other, the biological pattern may not match the cognitive phenotype, or the result may not fit the course described by the patient and informant. The correct response is to check what was actually measured, which method and reference interpretation the laboratory used, whether sample quality was acceptable, and whether the original clinical question remains appropriate. It is unsafe to “average” contradictory axes into a vaguely positive panel.

Similarly, a result close to a laboratory decision boundary should not be converted into categorical certainty by rhetoric. Reports may include interpretive language, but the requesting clinician remains responsible for integrating that language with the pre-test formulation. Where the result could materially change counselling or eligibility for a further intervention, discussion with the reporting laboratory or a relevant specialist is preferable to guessing.

PITFALL

Do not import a cut-off from a paper, another laboratory or memory into the current report. Fluid concentrations and ratios require the method-specific reference interpretation supplied by the laboratory. This fragment intentionally teaches direction, not an unverified universal threshold.

There is also a semantic trap in the word “positive.” A positive amyloid-axis classification may correspond to a low measured concentration in CSF. The classification and the raw direction are therefore not synonyms. State both clearly: the measured amyloid signal is reduced, and this

supports positivity of the relevant biological axis according to the laboratory's interpretive framework.

The CSF examination procedure, collection technique and generic constituents are not repeated here; they are covered in the Neurology foundations for psychiatrists notes. For this topic, the psychiatrist's responsibility is to formulate the dementia-specific question, understand the panel's axes, anticipate how the result will be used, and interpret the report without exceeding the evidence it contains.

22.5 Emerging blood biomarkers

Blood biomarkers are attractive because venous sampling is less invasive and potentially easier to repeat or deploy than CSF assessment. The emerging plasma panel includes **amyloid-beta-42/amyloid-beta-40 ratio, neurofilament light, total tau, phosphorylated tau-181 and possibly phosphorylated tau-217**. These candidates span amyloid, tau and neuronal-injury biology rather than representing a generic interchangeable blood test.

The wording “emerging” matters. It signals an evolving clinical role, not permission to assign every candidate the evidentiary status of an established diagnostic standard. In particular, the newest phosphorylated-tau candidate may be named, but the available evidence here does not provide its disease direction, a decision threshold, or diagnostic performance. An answer that supplies such figures from memory is more dangerous than an answer that explicitly states the limitation.

Conceptually, a blood test can occupy different places in a pathway. It may help decide who needs more definitive biological assessment, support a specialist formulation, or contribute to follow-up in a defined service. These are different uses, each requiring its own validation and clinical governance. Evidence for a use cannot automatically be transferred to another. Nor should “measurable in plasma” be paraphrased as “diagnostic in isolation.”

Adoption should therefore proceed from intended use to evidence, rather than from technological availability to indiscriminate ordering. A service needs to decide who is eligible for testing, what result categories mean, how uncertain findings are communicated, and which confirmatory route follows each category. Without that pathway, a convenient sample can generate an inconvenient ambiguity. The appropriate question is not simply whether a marker can be measured, but whether measuring it here will produce a clinically accountable decision.

- Identify the exact plasma analyte and the biological axis it is intended to represent.
- Ask whether the assay and interpretive pathway are validated for the clinical use being proposed.
- Preserve clinical confirmation and an escalation route for unexpected or discordant results.

The blood and CSF names can look deceptively similar. That does not make their numerical values, reference intervals or decision boundaries interchangeable across specimen types. The sample matrix, assay and intended use belong in the interpretation. A clinician should therefore

resist copying a CSF logic statement into a plasma report without checking what the blood assay has actually been validated to mean.

> **GOOD TO REMEMBER**
>
> When a report mentions an emerging plasma marker, document the clinical question, assay name, laboratory interpretation and intended next step. If the result is discordant with the phenotype or would drive a consequential decision, seek specialist or laboratory clarification rather than treating the flag as self-explanatory.

22.6 A practical reporting sequence

A disciplined report can be brief without being simplistic. Begin with the clinical syndrome and the reason biological clarification was sought. Name the specimen and measured analytes. Describe the amyloid, tau-pathology and injury axes separately. Then give an integrated interpretation, including discordance and limitations. Finish with the action the result supports, such as further assessment, review of an alternative explanation, or specialist discussion.

An adequate conclusion should remain traceable. The reader ought to be able to move backwards from the integrated sentence to the axis assignment, from the axis assignment to the measured analyte, and from the analyte to the laboratory's interpretation. Traceability makes disagreement productive: a colleague can identify whether the dispute concerns the clinical phenotype, the biological framework or the analytical result. A conclusion that merely says “biomarker positive” hides all of these distinctions.

Use precise language. “Compatible with” or “supports” is generally safer than wording that implies the laboratory has independently diagnosed the whole clinical disorder. Avoid describing total tau as a disease-specific marker, because its framework role is neuronal injury. Avoid describing phosphorylated tau merely as injury when the relevant framework assigns it to aggregated tau. Avoid calling the plasma panel established merely because its analytes have familiar names.

The pre-test and post-test conversations should mirror each other. Before sampling, explain the question and possible categories of result. After sampling, explain what biological process the result supports, what it does not establish, and whether it changes the working formulation. The patient and family need the clinical meaning, not a recital of abbreviations. Uncertainty should be stated plainly when the phenotype and biomarker pattern do not align.

Quality assurance begins before interpretation. Confirm identity and specimen type, use the reporting laboratory's own reference interpretation, and notice whether any analyte is missing or technically qualified. A partial panel should not be silently treated as a complete biological profile. If serial samples are being compared, the clinician should ensure that the comparison is methodologically meaningful rather than inferring biological change from unexamined laboratory differences.

Finally, keep the framework subordinate to the clinical record. Biomarkers can refine aetiology, but they do not quantify the person's everyday abilities, behavioural burden, safety, decision-

making needs or caregiver strain. Those require direct assessment. The highest-quality answer therefore links biology back to the person: what uncertainty was reduced, what uncertainty remains, and what action is justified now.

22.7 Examination synthesis

A strong short note has a simple spine. Define fluid biomarkers as biological adjuncts to clinical assessment. State the core CSF direction pattern. Explain that the amyloid concentration and ratio represent A, phosphorylated tau represents T, and total tau represents (N). Then name the emerging plasma candidates, including the newest phosphorylated-tau analyte, while explicitly refusing unsupported thresholds or performance claims.

For a long answer, add the distinction between syndrome and biology, explain why the axes must be read separately, describe how discordance is handled, and show how a laboratory report is integrated with phenotype and course. Mention that CSF procedural teaching and imaging physics sit in their respective chapters. This demonstrates ownership of dementia-specific biomarker interpretation without drifting into lumbar-puncture technique or modality science.

The conceptual hierarchy therefore begins with the clinical syndrome, proceeds through the biological axes and individual analytes, and reaches numerical interpretation only within the validated laboratory context. Keeping that order prevents both underuse and overclaiming. The candidate who remembers only “amyloid down, tau up” has the direction; the candidate who also separates T from (N), recognises plasma testing as emerging, and states the limits of the evidence has the complete answer.

23 · Cholinesterase inhibitors and memantine

Why this matters clinically. Cognitive enhancers are among the few dementia drugs that psychiatrists prescribe repeatedly across outpatient follow-up, liaison work and shared care. Examination answers therefore need more than a list of names. A competent answer links the transmitter target to the individual drug, matches the drug to the stage and type of dementia, gives a safe titration plan, anticipates adverse effects and explains how benefit will be judged when dramatic improvement is unlikely.

These medicines are symptomatic treatments. Their value is assessed through the patient's cognition, everyday function, behaviour relevant to care, tolerability and the caregiver's account of change. A small apparent gain on a test does not rescue a regimen that causes vomiting, syncope or loss of adherence; equally, absence of visible improvement does not by itself prove absence of benefit when the clinical aim may be preservation of function. Prescribing is therefore a structured therapeutic trial rather than an automatic response to a diagnostic label.

CHOLINERGIC	GLUTAMATE	TITRATE	REVIEW
AMPLIFY SIGNAL	LIMIT EXCESS	BY TOLERABILITY	FUNCTION AND SAFETY

23.1 Therapeutic logic and mechanisms

The **cholinergic hypothesis** proposes that progressive loss of cholinergic neurons and declining acetylcholine contribute to cognitive deterioration in Alzheimer dementia. Acetylcholinesterase

and butyrylcholinesterase both degrade acetylcholine. Inhibiting this degradation prolongs the availability of acetylcholine at functioning synapses. This is a rational symptomatic strategy, but cholinergic depletion is not considered the sole explanation of Alzheimer disease. The mechanism therefore supports modest therapeutic expectations, not a claim that the underlying neurodegenerative process has been removed.

Donepezil is a selective, reversible acetylcholinesterase inhibitor. **Rivastigmine** inhibits both acetylcholinesterase and butyrylcholinesterase and is described as a reversible, non-competitive inhibitor. **Galantamine** is a competitive, reversible and selective acetylcholinesterase inhibitor with additional nicotinic-receptor agonist properties. These distinctions are examinable because they explain why the drugs belong to the same therapeutic family without being pharmacologically identical.

Memantine acts through a different transmitter system. It is a low-to-moderate-affinity, non-competitive antagonist at N-methyl-D-aspartate receptors. Preferential binding to open receptor-operated calcium channels makes the block activity dependent: pathological glutamatergic activation and excessive ion flux are restrained while ordinary signalling is less indiscriminately suppressed than it would be by a high-affinity, persistent block. The therapeutic concept is attenuation of excitotoxic glutamate activity, not enhancement of acetylcholine.

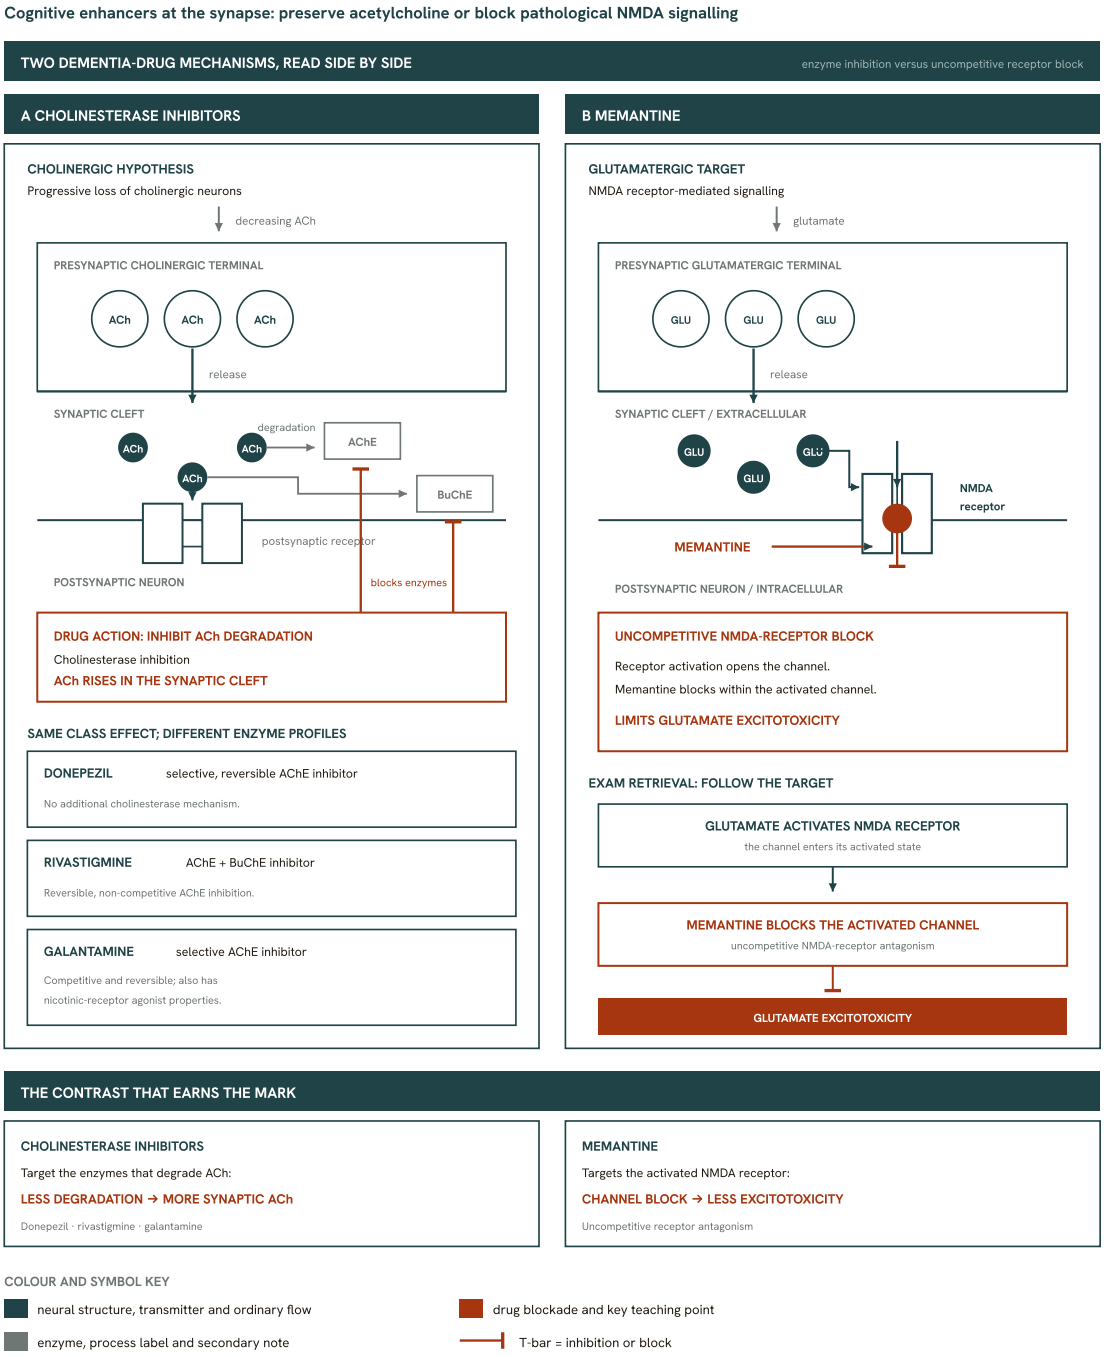


Fig 32.10 **Cognitive enhancers at the synapse.** Cholinesterase inhibitors raise synaptic acetylcholine by preventing its degradation: donepezil selectively and reversibly inhibits AChE; rivastigmine inhibits AChE and BuChE; galantamine competitively and reversibly inhibits AChE and also has nicotinic-receptor agonist properties. Memantine instead produces uncompetitive NMDA-receptor block, limiting glutamate excitotoxicity. (Maudsley Prescribing Guidelines in Psychiatry, 15th ed.)

The accompanying figure should be read as parallel symptomatic strategies at the synapse. Cholinesterase inhibition increases persistence of available acetylcholine; memantine modifies excessive receptor-mediated glutamatergic flux. The different mechanisms also explain why concurrent use can be rational in the appropriate stage of Alzheimer dementia.

AGENT	PRIMARY PHARMACOLOGICAL ACTION	CLINICAL IMPLICATION
Donepezil	Selective reversible acetylcholinesterase inhibition	Oral cholinergic strategy with simple daily administration
Rivastigmine	Acetylcholinesterase plus butyrylcholinesterase inhibition	Oral and transdermal routes allow route-based individualisation
Galantamine	Selective reversible acetylcholinesterase inhibition plus nicotinic activity	Oral cholinergic strategy in extended-release or solution form
Memantine	Activity-dependent N-methyl-D-aspartate receptor antagonism	Non-cholinergic option for later-stage Alzheimer dementia and combination treatment

RULE **Mechanism determines the therapeutic family; stage, tolerability and route determine the prescription.** Do not choose among these agents merely by remembering that all are called cognitive enhancers.

23.2 Indications and choosing an initial strategy

The **cholinesterase-inhibitor indication envelope** includes donepezil, rivastigmine and galantamine for mild-to-moderate dementia due to Alzheimer disease, with recommendations also supporting their use in severe Alzheimer disease. Rivastigmine additionally has a licence in some countries for mild-to-moderate dementia associated with Parkinson disease. Jurisdiction matters: a country-specific licence cannot be silently converted into a universal indication.

Within the shared Alzheimer indication, there is no sound clinical reason to present a cholinesterase inhibitor as universally best. Selection is practical and patient centred. Consider whether a daily oral regimen is feasible, whether a caregiver can supervise administration, whether swallowing or gastrointestinal tolerability makes a transdermal route attractive, whether the current medicine burden threatens adherence and whether cardiac history changes the safety threshold. The decision should be documented as an individual therapeutic trial with an agreed method of review.

A **sustained-release high-dose donepezil formulation** has United States regulatory approval for moderate-to-severe Alzheimer disease, carries a greater early gastrointestinal adverse-effect burden than the usual formulation and is not approved in the United Kingdom. The important lesson is not to import a formulation or titration practice from another jurisdiction without checking current country-specific product information. Higher dose is not a synonym for better individual treatment.

The **memantine indication** is moderate-to-severe dementia due to Alzheimer disease. NICE positions it as an option in moderate Alzheimer disease when acetylcholinesterase inhibitors are contraindicated or not tolerated, and as treatment in severe Alzheimer disease. Thus memantine may be an alternative when a cholinergic strategy cannot be used, or part of a later-stage regimen rather than a routine substitute at every stage.

TRAP **Do not write that cholinesterase inhibitors and memantine are interchangeable.** They differ in mechanism and stage-related indication. Memantine is not simply the "stronger" cognitive enhancer, and severe disease does not automatically require stopping a tolerated cholinesterase inhibitor.

23.3 Initiation and titration

Titration is part of treatment, not administrative detail. Before starting, establish what the caregiver will observe, record relevant baseline function, review pulse and cardiovascular history, examine the full medicine list, and decide who will supervise dosing. Avoid simultaneous changes to cognitive enhancers where possible. This makes emerging nausea, dizziness, insomnia or syncope interpretable and prevents disease progression from being confused with a treatment effect.

DRUG AND DEMENTIA INDICATION	STARTING, USUAL AND TITRATION REGIMEN
Donepezil for Alzheimer dementia	**Start 5 mg once daily; usual dose 10 mg once daily; increase by 5 mg daily only after at least 4 weeks.**
Rivastigmine for Alzheimer dementia	**Start oral treatment at 1.5 mg twice daily, or use the 4.6 mg per 24 hours patch. The usual oral dose is 3-6 mg twice daily, increased by 1.5 mg twice daily at intervals of at least 2 weeks; the usual patch is 9.5 mg per 24 hours after at least 4 weeks. A 13.3 mg per 24 hours patch may be considered after 6 months. An oral dose above 6 mg daily may be switched directly to the 9.5 mg per 24 hours patch.**
Galantamine for Alzheimer dementia	**Start extended-release treatment at 8 mg once daily, or solution at 4 mg twice daily. The usual extended-release dose is 16-24 mg once daily, increased by 8 mg daily after at least 4 weeks; the usual solution dose is 8-12 mg twice daily, increased by 4 mg twice daily after at least 4 weeks.**
Memantine for moderate-to-severe Alzheimer dementia	**Start 5 mg once daily; the usual dose is 20 mg daily, either once daily or as 10 mg twice daily; increase by 5 mg at intervals of at least 1 week.**

The table gives the regimen, but good prescribing adds method. Explain the formulation clearly, especially when oral solution or a patch is used. Confirm that the caregiver understands which previous formulation is being replaced. At every increase, ask whether the previous dose was actually tolerated and taken. A nominal schedule should never force escalation through clinically important adverse effects.

A formulation switch should be treated as a prescribing event, not as clerical substitution. Reconcile the previous oral or patch regimen, write the replacement route unambiguously and verify that the superseded supply will not continue by mistake. At follow-up, ask how the medicine is actually given rather than asking only whether the patient is "on" it. This detects

duplicate administration, missed doses, misunderstanding of solution volumes and practical failure of caregiver supervision without assuming deliberate non-adherence.

- Define the treatment target in observable terms, including everyday tasks and the level of prompting required.
- Record the formulation, route, supervisor and intended review point.
- At review, separate adherence failure, adverse effects, intercurrent illness and progression of dementia.
- If benefit and harm are uncertain, reconstruct the timeline with the caregiver before changing several medicines.

GOOD TO REMEMBER

The caregiver is part of the pharmacological intervention. A technically correct prescription fails if nobody can describe how it is administered, what changed after titration, or which new symptom should trigger review.

23.4 Cholinergic adverse effects and their interpretation

Cholinergic adverse effects commonly include nausea, vomiting, diarrhoea, dizziness and insomnia; urinary incontinence has also been reported. They are most likely near treatment initiation or after a dose increase, are dose related and are often transient. This temporal pattern is clinically useful. A symptom that appears after escalation should initially be treated as a possible drug effect rather than automatically labelled deterioration of dementia.

Management begins by grading severity and consequence. Mild transient nausea in a patient who remains hydrated and adherent poses a different problem from persistent vomiting, postural instability or diarrhoea causing frailty. Review when the symptom began, which formulation was used, whether a dosing error occurred and whether another medicine or illness offers a better explanation. The response may be to pause escalation, return to a previously tolerated regimen, reconsider route or reconsider the drug. The purpose is not to “push through” toxicity to reach a table dose.

MNEMONIC

Gut, sleep, bladder, beat

For cholinesterase inhibitors, remember gastrointestinal symptoms, insomnia, urinary incontinence and slowing of the cardiac beat. Dizziness may be the presenting clue that links the autonomic and cardiac consequences.

Adverse effects can masquerade as the illness being treated. Insomnia may worsen daytime cognition; dizziness may reduce mobility; urinary incontinence may be misread as loss of functional stage; and vomiting may produce apparent global decline. The correct response is a medication timeline and physical assessment, not reflex escalation of dementia treatment.

PITFALL

Do not attribute new vomiting, syncope or falls to “advanced dementia” without checking the recent cholinesterase-inhibitor start or dose increase. That shortcut can convert a reversible adverse effect into injury and further disability.

23.5 Cardiovascular safety and the patient at higher risk

The clinically important vagotonic effect is slowing of heart rate. **Bradycardia and syncope** are the most prominent cardiovascular adverse effects of cholinesterase inhibitors. Their consequences extend beyond a pulse reading: syncope may lead to falls or fractures, and clinically significant bradyarrhythmia may result in pacemaker placement. Patients with sick sinus syndrome or conduction disturbance require particular caution. People with dementia with Lewy bodies may be especially susceptible to bradyarrhythmic effects.

Cardiovascular screening should be proportionate rather than ritualistic. Take a history of syncope, unexplained falls, palpitations and known conduction disease; examine the pulse; and review drugs that also slow heart rate. An electrocardiogram is advised when cardiovascular disease is present or when concomitant negative-chronotropic medicines are used. If syncope develops, assess it as a potentially consequential cardiac event rather than merely recording “dizziness”.

Risk assessment continues after initiation. A normal baseline assessment does not guarantee future tolerance after dose escalation or addition of another rate-slowing drug. Conversely, a risk factor is not an automatic prohibition. It changes the threshold for investigation, shared decision-making, titration and review. The practical question is whether the expected cognitive or functional value justifies the cardiovascular uncertainty for this patient.

The signal around galantamine in mild cognitive impairment is a separate safety lesson, not an Alzheimer-dementia mortality estimate. In trials involving mild cognitive impairment, mortality was **1.5% with galantamine versus 0.5% with placebo**, with cardiovascular deaths contributing substantially. This led to a United States regulatory warning restricting galantamine use in mild cognitive impairment; the relevance of that signal to Alzheimer dementia remains unclear. The examination error is to transfer the result to a different population without preserving its qualifier.

23.6 Memantine tolerability and combination treatment

Memantine tolerability is generally favourable, but adverse effects still require active inquiry. Reported common effects include somnolence, dizziness, balance disturbance, constipation, headache, hypertension, dyspnoea and elevated liver-function values. Warnings particularly concern severe hepatic impairment and epilepsy or seizures. When a patient becomes more sleepy, unsteady or constipated after initiation, the change should not be absorbed into a vague narrative of dementia progression.

The adverse-effect profile differs from cholinergic toxicity, which helps when memantine is used because a cholinesterase inhibitor was not tolerated. It does not make memantine risk free.

Balance disturbance and dizziness may be functionally important even when the drug is otherwise easy to take. The review should ask what the symptom costs in mobility, supervision and safety, not merely whether it appears on a checklist.

The **cholinesterase-inhibitor plus memantine combination** is recommended by NICE over cholinesterase-inhibitor monotherapy in moderate-to-severe Alzheimer disease: addition may be considered in moderate disease and should be offered in severe disease. There are no pharmacokinetic or pharmacodynamic interactions between these classes. This permits rational concurrent treatment, but it does not eliminate the need to identify which agent is associated with a new adverse effect.

Combination treatment should retain a clear rationale. The patient should have a stage-appropriate indication, a tolerable existing regimen and a practical plan for observing change after the added medicine. If several drugs are changed together, later attribution becomes difficult. The absence of a class interaction means co-prescribing is pharmacologically permissible; it does not mean that every patient will derive enough additional value to justify indefinite treatment.

23.7 Longitudinal prescribing: LOCUS, FOCUS and MODUS

LOCUS. Initiation and routine review usually belong in organised outpatient or memory-service care with a reliable caregiver and access to physical monitoring. Acute assessment becomes appropriate when there is syncope, clinically important bradycardia, repeated vomiting, marked dehydration, sudden loss of balance or another change that cannot safely wait for routine follow-up. The care setting should follow the immediate risk, not the prestige of the clinic that started treatment.

FOCUS. Early treatment focuses on tolerability, correct administration and a trustworthy baseline. Short-term review focuses on adverse effects and whether titration remains justified. Later review focuses on cognition in context, everyday function, caregiver-observed change, falls, adherence and evolving treatment goals. A score can contribute evidence, but it cannot independently decide continuation or withdrawal.

MODUS. Drug treatment sits inside a broader care plan. Prescribing work includes explanation, caregiver training, formulation choice, medication reconciliation, physical monitoring and coordination with the clinician managing cardiovascular or neurological comorbidity. The relevant guidance should be checked before prescribing: current country-specific product information governs available formulations and licensing, while NICE guidance supplies a clear stage-based framework for memantine and for combination treatment. Local institutional policy should determine practical monitoring and responsibility for follow-up.

A useful review reconstructs a sequence: what the patient could do before treatment, what changed after initiation, what changed after each increase, which adverse effects appeared, and what happened to caregiver burden and supervision. This is more informative than a binary question about whether memory “improved”. It also protects against opposing errors: stopping a

tolerated drug solely because dementia continues to progress, and continuing a harmful drug solely because deterioration is expected from the disease.

Continuation and withdrawal deserve the same discipline as initiation. Review the original indication, current stage, adherence, adverse effects and the functional observations that matter to the patient and caregiver. Avoid treating progression alone as proof that the medicine has no value, because the counterfactual course is not directly observable. Equally, avoid indefinite continuation when administration has become burdensome, harms are prominent or the regimen no longer serves the agreed goals. Where change is planned, communicate what will be observed and who will respond if clinically important deterioration or relief of adverse effects follows.

Shared decisions should acknowledge uncertainty without becoming vague. State the intended benefit, the harms that would prompt early contact, the method of administration and who will observe the outcome. When tolerability is poor, reassess dose, route and drug choice rather than assuming the class has failed. When treatment seems ineffective, verify adherence and intercurrent illness before drawing a conclusion. When goals of care change, reconsider whether the current regimen still serves the patient's function and comfort.

> The must-memorise principle: choose by indication and route, titrate only after tolerability is established, and judge treatment by function and safety as well as cognition.

24 · Behavioural and psychological symptoms of dementia: assessment, management and antipsychotic risk

The marks lie in the sequence, not in a drug list. The candidate must recognise the syndrome, search for what the person may be communicating, formulate the interaction between illness, body, environment and care, and use medication only when the target and the risk justify it. **Behavioural and psychological symptoms of dementia** include aggression, agitation, vocalisation, distress during care, disinhibition, hallucinations, delusions, apathy, low mood and anxiety. They are common across the course of dementia and vary greatly in form, severity and meaning.

These symptoms occur to some degree in **over 90% of people with dementia**. That prevalence should not turn them into an automatic indication for sedation. The central clinical task is to translate a broad label into an observable target, a plausible formulation and a proportionate plan. Management is anchored in Maudsley prescribing guidance for person-centred non-drug care and in the Kaplan and Sadock risk synthesis when antipsychotics are considered.

ASSESS	FORMULATE	REVIEW
DEFINE THE TARGET	FIND THE DRIVER	CONTINUE ONLY WITH BENEFIT

24.1 From a label to a target syndrome

BPSD is an umbrella description, not a unitary diagnosis. “Agitation” may refer to pacing, repeated calling, resistance during bathing or striking a caregiver. “Psychosis” may refer to a fixed false belief, a perceptual experience, a misinterpretation of the surroundings or frightened speech

that has been incompletely elicited. Apathy, low mood and anxiety can look superficially similar because each may present as withdrawal. The assessment must therefore move from the caregiver's summary to a description of what actually happens.

Characterise the behaviour in neutral language: its form, context, apparent trigger, consequence, variability and impact. Establish whose distress is being treated. A behaviour may endanger the person or another individual; it may instead be inconvenient to a rigid routine without being harmful. This distinction matters because the burden of drug adverse effects falls on the person with dementia, even when the immediate request for treatment comes from others.

RULE **Name the target before naming the medicine.** "BPSD" is too broad for a treatment trial; the prescription must be tied to a specific, observable and consequential symptom.

The same outward behaviour may have different meanings in different people. Repeated attempts to leave may express a familiar life routine, fear, discomfort or a need for contact. Resistance during care may be an understandable response to unfamiliar touch or rushed communication. This is why phenomenology and biography sit beside the physical examination rather than after it.

24.2 Structured assessment: body, context and sequence

Begin with a safety and medical screen. Look for pain, infection, constipation and adverse effects of medication before concluding that the dementia itself has worsened. Review sensory difficulty, mobility and falls vulnerability. Observe the person directly when possible, because caregiver language can compress several different events into a single term. If the change appears acute, delirium must be considered; its detailed assessment belongs in the Stroke, TBI, delirium and focal brain syndromes notes.

Next, reconstruct the event. Record what happened before the behaviour, the behaviour itself and what followed. An antecedent-behaviour-consequence chart is useful because consequences may inadvertently maintain a pattern: the person may escape an overwhelming task, gain reassurance or encounter escalating confrontation. The chart is not a blame instrument. It turns impressions into testable observations and permits comparison after a change in care.

OBSERVED PRESENTATION	QUESTIONS THAT SHARPEN THE TARGET	USEFUL FORMULATION OUTPUT
Agitation, aggression or vocalisation	What precedes it, who is present, and what ends it?	Trigger, function, risk and a measurable behavioural target
Distress during care	Which task, approach, touch or pace is associated with distress?	A modified care interaction rather than a reflex prescription
Hallucinations or delusions	What exactly is experienced or believed, and how distressing is it?	Phenomenology, associated risk and need for treatment
Apathy, low mood or anxiety	Is there loss of initiative, subjective distress, fear or avoidance?	A discriminated emotional or motivational target
Disinhibition	In which settings does it occur and what consequences follow?	Environmental, supervisory and behavioural safeguards
Variable sleep or pain behaviour	What pattern appears across observation and care records?	A temporal clue to a physical or environmental driver

Life history adds meaning to the sequence. Familiar routines, preferences, occupation, language and prior responses to stress can explain why an apparently irrational event is coherent from the person's perspective. Sleep and pain records, direct observation and a shared formulation meeting convert these data into a care plan that caregivers can apply and review. This unmet-need formulation reframes "problem behaviour" as an expression of distress until evidence suggests otherwise.

24.3 Formulation and risk: deciding what requires action

A useful formulation has predisposing vulnerability, immediate precipitants, perpetuating responses and protective resources, but it remains practical. It should answer what is happening, why it may be happening now, what harm could follow, and what change can be tested. The output is not merely a diagnosis. It is a shared hypothesis linking an intervention to an observable outcome.

Risk assessment separates urgency from annoyance. Consider danger to the person, danger to others, interruption of essential care, inability to meet basic needs and the cumulative effect of sustained distress. Also examine how staff or family responses alter the event. Fear, argument, rapid commands and inconsistent limits can intensify an interaction without anyone intending harm. Conversely, a calm approach, adequate time and predictable sequencing may allow necessary care without medication.

MNEMONIC

PAIN before prescription

Physical illness and pain; **A**dverse medication effects; **I**mpaired sensory input; **N**eeds, narrative and the care context. It is a compact prompt for formulation, not a substitute for examination or observation.

Document the baseline in terms the team can recognise. “Less agitated” is weak; a description of the relevant care episode, its severity and its consequences is clinically usable. Agree in advance what improvement would count as meaningful, which adverse outcomes will prompt reconsideration, and who will review the plan. This prevents an intervention from continuing simply because it has become part of the medication chart.

24.4 Management architecture: LOCUS, FOCUS and MODUS

LOCUS asks where safe care can occur. Outpatient or home-based care is appropriate when risk can be contained and caregivers can implement the plan. Inpatient care may be needed when assessment, containment or treatment cannot be delivered safely in the usual setting. An emergency presentation requires immediate protection and a search for reversible physical or environmental drivers, followed by reassessment once the crisis has settled.

FOCUS changes with phase. In an acute crisis the target is immediate safety and relief of severe distress. Over the short term, the task is to identify precipitants and make the care environment workable. In the middle term, behavioural learning, caregiver consistency and restoration of tolerable routines become central. Long-term care aims to preserve function and dignity while repeatedly removing interventions that no longer have a clear benefit.

MODUS means that management is genuinely multimodal: treatment of physical contributors, communication and environmental modification, behavioural intervention, caregiver psychoeducation, selected psychological or sensory approaches, social support and, only when justified, medication. The evidence favours **first-line personalised multicomponent non-drug care** delivered with caregivers. Drug evidence is less secure and several agents carry serious adverse effects.

- Define the priority target and the harm or distress attached to it.
- Correct contributory physical and sensory factors before escalating treatment.
- Choose an intervention that follows directly from the formulation.
- Teach caregivers how to apply the plan consistently across settings.
- Record response and adverse effects, then revise rather than merely add treatment.

24.5 Non-pharmacological treatment in practice

At this management stage, **LOCUS** determines feasibility: a home plan must fit the household, while a ward plan must survive shift changes. **FOCUS** is the target behaviour and the unmet need inferred from assessment. **MODUS** begins with changing the interaction or environment, not trying to suppress the person. Simplify demands, pace care, use familiar routines, reduce avoidable confrontation and build pleasant engagement around the person’s history and preferences.

Behavioural management and caregiver psychoeducation are generally successful, and their effects can persist for months. Their strength lies in repeated application: caregivers learn to recognise antecedents, alter their approach and notice early escalation. A written care plan should

specify what helps, what worsens the situation and how to respond consistently. Review remains necessary because the person's cognition, health and environment change.

Among sensory-stimulation approaches, **music therapy** has the most convincing evidence for reducing agitation and aggressive behaviour. Multicomponent interventions, animal-assisted therapy, doll therapy, bright light therapy and aerobic exercise may help. Aromatherapy has not shown benefit. These findings should guide selection without turning any modality into a ritual. An intervention is useful when it suits the person, is deliverable in the actual locus of care and changes the agreed target.

Pain assessment and effective treatment deserve special attention because pain may be expressed behaviourally when self-report is limited. Even without overt pain, an analgesic trial, usually with paracetamol, may be worthwhile. The trial still requires a target and review: improved comfort during movement or care is more informative than indefinite continuation without observed benefit.

> **GOOD TO REMEMBER**
>
> Ask the caregiver to show you the difficult interaction or describe it step by step. A direct account often reveals a modifiable trigger that disappears inside the single word "agitation".

24.6 When medication enters the plan

Medication is not the next rung after reassurance fails. In the relevant **LOCUS**, establish beforehand that physical illness, pain, constipation, medication effects and sensory deficits have been addressed. The **FOCUS** must be severe distress, serious risk of physical harm or another clearly consequential target, rather than wandering, calling out or care resistance considered in isolation. The **MODUS** remains combined: a drug trial supplements a person-centred plan and does not replace it.

Antipsychotic use for BPSD is discouraged as routine practice because average benefit is small, tolerability is poor and mortality is increased. Conventional agents other than haloperidol have not shown efficacy in the evidence summarised here. Among newer agents, risperidone has evidence for psychosis, agitation and overall BPSD; olanzapine for agitation; and aripiprazole for overall BPSD. Quetiapine has generally failed to show effectiveness except at doses that may be poorly tolerated. These are evidence signals, not instructions to cycle through agents.

Before prescribing, discuss the expected target, uncertainty of benefit, material risks and the intended stopping plan with the caregiver. Assess falls vulnerability and stroke risk factors. Record why the situation crosses the medication threshold and why less hazardous measures are insufficient. If the presentation does not involve severe distress or serious risk of physical harm, neither risperidone nor olanzapine should be used merely to make care easier.

TRAP **"Agitation in dementia" is not itself an antipsychotic indication.** The high-scoring answer states the target, excludes physical and environmental drivers, documents non-drug measures, and then explains why risk or severe distress makes a time-limited trial proportionate.

24.7 Agent selection and dementia-scoped dosing

For **persistent aggression in moderate to severe Alzheimer dementia** that has not responded to non-pharmacological approaches and carries risk of harm, risperidone is the preferred licensed option in the referenced UK guidance. Its indication is short term, **up to 6 weeks**. Haloperidol is also licensed there, but its dangers make risperidone the preferred agent. Licensing is jurisdiction-specific and should never be silently transferred to Indian practice.

Dosing must remain indication-bound. The optimal risperidone dose described for BPSD in dementia is **500 micrograms twice daily (1 mg/day); it is licensed up to 1 mg twice daily**. This is deliberately a dementia-scoped low-dose approach. Schizophrenia-range doses do not belong in this section. The broader teaching of antipsychotic pharmacology and extrapyramidal and metabolic adverse effects belongs in the Schizophrenia: management, antipsychotics and adverse effects notes.

The same indication discipline applies if another antipsychotic is considered: use dementia-specific low-dose ranges and do not import a dose from a different disorder. Choice should be driven by the target for which an agent has evidence, the person's vulnerabilities, likely tolerability and the capacity to review promptly. A long medication list is not sophistication. It usually means that targets and stopping decisions have become blurred.

Between-agent mortality comparisons do not provide a simple "safe antipsychotic". Haloperidol has generally been associated with greater mortality, while quetiapine users showed lower risk in some comparative evidence. Clinically meaningful differences were not found for several other agents, and another comparative analysis did not find significant differences among commonly used newer antipsychotics. Risk appeared strongest soon after initiation, and higher dose was generally associated with greater mortality except in the quetiapine findings. These observations support caution, not substitution by reputation.

24.8 Boxed warning, cerebrovascular risk and review

The decisive safety fact is the **antipsychotic boxed warning in dementia**. Placebo-controlled evidence underlying the warning found a **1.6 to 1.7-fold increase in all-cause mortality versus placebo** among people with dementia taking atypical antipsychotics. Deaths were attributed to cardiac events, including heart failure and sudden death, or to infections, most often pneumonia. The benefit for aggression or psychosis is modest and short term, while the adverse outcome is potentially fatal.

Product labelling also warns of possible **cerebrovascular events** when antipsychotics are used in dementia, and the warnings now apply across the class. Do not attach a remembered numerical stroke estimate to this statement: the permitted evidence gives a quantitative mortality estimate but no quantitative relative risk specifically for stroke. Mortality and cerebrovascular risk are related safety concerns, not interchangeable numbers.

In any **LOCUS**, the **FOCUS** of review is benefit against the original target, emerging harm and whether the indication still exists. The **MODUS** is active deprescribing, not passive renewal. Use the **lowest effective dose with repeated stopping review**. Review at **4 to 6 weeks, and taper**

or withdraw when there has been no significant response after 4 weeks. Consider stopping at every review even when some benefit is reported.

> **PITFALL**
> A drug started during an inpatient crisis may remain on the discharge prescription after the trigger has resolved. The transition of care must carry the original target, observed response, adverse effects and a stopping plan, not merely the medicine name.

24.9 Other drug options and the final hierarchy

Other medicines do not remove the need for formulation. In the chosen **LOCUS**, set a symptom-specific **FOCUS** and retain psychosocial measures as the surrounding **MODUS**. **Non-antipsychotic drug options** have uneven evidence. Valproate trials for dementia-related behaviour have been disappointing. Selective serotonin reuptake inhibitors are most helpful for agitation, with citalopram having the strongest evidence, but its cardiac repolarisation risk constrains use in older people. Benzodiazepines are poorly supported and are associated with cognitive decline, pneumonia, all-cause mortality, falls and hip fractures.

Brexpiprazole, a dopamine partial agonist, improved agitation in Alzheimer dementia compared with placebo and is the only medicine approved by the United States regulator specifically for agitation associated with dementia due to Alzheimer disease. It is not available in the UK. Regulatory approval does not cancel class risk, replace non-pharmacological care or convert every episode of agitation into a prescribing indication. Availability and approval in India require current local verification.

The management hierarchy is therefore coherent. Begin by defining the behaviour and the person's distress. Then treat bodily contributors, understand the sequence, change care and environment, and work with caregivers. If a medicine becomes proportionate, choose it for a named target, use a dementia-scoped low dose, explain mortality and cerebrovascular warnings, and make withdrawal part of the original prescription. A careful decision not to prescribe is active treatment when the formulation points elsewhere.

> In BPSD, personalised non-drug care is the default; an antipsychotic is a low-dose, target-bound, time-limited trial whose mortality and cerebrovascular risks must be explained and repeatedly reviewed.

25 · Caregiver support, psychosocial interventions and end-of-life care in dementia

Dementia care is delivered through a relationship, not to an isolated brain. The person with dementia remains the centre of care, but the caregiver commonly becomes the observer, interpreter, organiser and practical carrier of the treatment plan. **Caregiver stress, isolation and burden** may include psychological, physical and financial problems. Caring for cognitive impairment also creates greater psychological strain than caring for medical illness, even though the caregiver remains the key resource through whom most interventions are implemented. This combination of indispensability and vulnerability defines the clinical task.

The examination answer therefore needs more than an instruction to “counsel the family”. It should show how the clinician assesses caregiver burden, builds an active support package, uses the caregiver to make psychosocial treatment workable, and changes the goal of care when advanced disease makes comfort and dignity paramount. The strongest answer treats support as part of treatment architecture. It also keeps the boundary clear: detailed assessment and prescribing for behavioural and psychological symptoms belong in the corresponding management section, while this section concentrates on the caregiver and psychosocial system around the patient.

NOTICE	EQUIP	SUSTAIN	COMFORT
BURDEN AND RESOURCES	KNOWLEDGE AND SKILLS	SUPPORT AND RESPITE	DIGNITY AT THE END

25.1 The caregiver as co-recipient and implementation partner

A caregiver is simultaneously a person with legitimate support needs and the person expected to translate recommendations into daily practice. When the support need is overlooked, the consultation becomes extractive: information and labour are requested from someone whose capacity to continue is never assessed. When implementation is overlooked, support may remain sympathetic but disconnected from the practical demands of care. A sound formulation holds these positions together.

The caregiver’s account supplies history that a brief encounter cannot. It is also shaped by fatigue, fear, relationship history and competing duties. This does not make it unreliable. The clinician must listen for patient information as well as caregiver distress. “Nothing works” may describe a care problem, an exhausted implementation system, or their interaction. Clarify what is happening, what has been tried, what is realistic, and what support would make the plan possible.

RULE **Link the patient outcome to the caregiver outcome.** Ask what should improve for the person with dementia and what must become more manageable for the caregiver who will carry the plan.

The phrase **caregiver as the key intervention resource** should never be mistaken for unlimited availability. Capacity fluctuates with sleep, health, employment, finances, family conflict and the amount of help available. A technically elegant intervention fails if it assumes time, privacy, physical strength or cooperation that the household does not possess. Feasibility is therefore a clinical property of the intervention, not a moral test of the family.

25.2 Assessing burden without reducing it to a score

Assessment begins by identifying who actually provides care. The named relative in the record may not be the person supervising meals, accompanying consultations or managing difficult evenings. Map the tasks, the people carrying them, the points at which responsibility changes hands, and the gaps where nobody is reliably available. Ask separately about emotional strain, physical demands, social isolation and financial pressure because a global question about “coping” often invites a socially acceptable answer.

Then examine the relationship between burden and the care situation. Which task is hardest? Which period feels least manageable? What behaviour evokes fear, anger, shame or helplessness? What has the caregiver stopped doing? What support is offered but unusable, and what support is absent? The aim is not to grade devotion. It is to identify the mechanisms by which strain is produced and the resources that may interrupt them.

ASSESSMENT DOMAIN	CLINICAL ENQUIRY	PLAN GENERATED
Care tasks	What must be done, by whom, and at which difficult points?	A realistic map of labour and supervision
Psychological strain	What emotions, thoughts or interactions are becoming hard to bear?	Support directed at the actual source of distress
Physical impact	How is caregiving affecting rest, health and bodily capacity?	Protection of the caregiver's ability to continue safely
Isolation	Which relationships, roles or activities have narrowed or disappeared?	Reconnection with usable social support
Financial pressure	Which care demands create direct or indirect economic difficulty?	A plan that does not assume unavailable resources
Knowledge and confidence	What is poorly understood, feared or difficult to perform?	Targeted education, rehearsal and feedback

Assessment should also identify strengths. These may include a calm communication style, an effective routine, a cooperative relative, a meaningful relationship or prior success with a behavioural strategy. Strengths are not decorative positives added after listing problems. They show where an intervention can attach. The clinician can reinforce what already works, reduce unnecessary complexity and reserve scarce caregiver effort for the change most likely to matter.

Revisit burden when care demands, health or living arrangements change. A feasible plan may become exhausting, while competence may grow with practice. Reassessment prevents the chart from freezing the caregiver into a permanent label such as "coping well" or "poor support".

25.3 What makes a caregiver intervention effective

Structured caregiver intervention is active treatment, not a loose collection of reassuring remarks. Group approaches work best when techniques are structured and intensive, caregivers participate actively, and the programme combines education, support and respite with knowledge transfer, skill building and competency. Individual support strengthens this package. Each component answers a different failure point in the care system, so merely naming them is not enough.

- **Education** gives the caregiver a coherent explanation of the care problem and the purpose of the plan.

- **Active participation** turns the caregiver from a passive recipient of advice into a collaborator who can test and refine it.
- **Skill building** converts general principles into observable actions that can be rehearsed and corrected.
- **Competency** means the caregiver can apply the skill in the real situation, not merely repeat the teaching point.
- **Support** creates space for distress, ambivalence and relationship strain without treating these reactions as failure.
- **Respite** interrupts continuous demand and helps preserve the person who is sustaining care.

The clinical sequence is explain, demonstrate, rehearse, adapt and review. Explanation without practice often produces apparent agreement but little change at home. Practice without explanation can feel controlling and may be abandoned when the situation varies. Review closes the loop by asking what the caregiver tried, what happened, what became easier, and where the technique broke down. The answer may require simplifying the skill, changing the environment or redistributing a task rather than repeating the same advice more forcefully.

MNEMONIC

CARE

Clarify the burden and target; **A**ctively involve the caregiver; **R**ehearse skills and arrange respite; **E**valuate competence, feasibility and support. The mnemonic keeps the intervention active and reviewable.

Group caregiver work offers shared learning and support, but it is not the whole intervention. A general strategy may remain difficult in a particular relationship or emotionally charged task. **Individual support** permits formulation at that level and discussion of shame, resentment, grief or conflict that may not be voiced in a group. The formats are complementary when linked by common goals.

25.4 Psychoeducation that changes behaviour

Psychoeducation should reduce confusion and increase usable agency. It begins with the caregiver's current model: what does the caregiver think the patient is doing, why is it happening, and what response seems sensible? If the caregiver sees every failure as deliberate refusal, repeated confrontation may follow. A more useful formulation can change the response without dismissing the caregiver's frustration. The teaching must remain specific to the person and the difficult situation.

Effective teaching uses plain explanation, an actionable target, demonstration and feedback. Ask the caregiver to describe the plan back and anticipate obstacles. Written instructions should identify what to do, what to avoid and how to judge benefit. "Maintain routine" is too abstract unless the caregiver knows which routine and how to adapt it.

Knowledge transfer, skill building and competency form a progression. Knowledge answers why. Skill answers how. Competency answers whether the caregiver can carry out the action under real conditions. This progression protects against a common clinical illusion: information was delivered, therefore care was improved. The meaningful endpoint is not the amount said in the consultation but the change the household can sustain.

> **GOOD TO REMEMBER**
> Before adding advice, ask the caregiver to take you through the hardest care episode from beginning to end. The point where the interaction becomes unmanageable usually gives a more useful teaching target than a broad request for "stress management".

Respite should be discussed as a treatment component rather than as permission to abandon responsibility. Caregivers may refuse help because the available arrangement feels unsafe, the patient rejects unfamiliar care, family expectations are rigid, or organising relief creates more work than it removes. The clinician should explore these barriers without moralising. A respite plan becomes real only when the caregiver can identify what demand will pause and who can safely hold it.

25.5 Caregiver-supported psychosocial care for the person with dementia

Psychosocial care becomes effective when it is personalised, multicomponent and embedded in the patient's ordinary environment. For behavioural and psychological symptoms, this caregiver-centred non-drug approach occupies the **first-line position**. The caregiver helps identify the target behaviour, the context in which it occurs, and the responses that may calm, worsen or maintain it. The clinician then converts this information into a plan that can be applied consistently.

Behavioural management with caregiver psychoeducation is centred on the individual patient's behaviour. Its logic is formulation rather than suppression: define what happens, observe the setting and sequence, modify the interaction or environment, and review the result. A personalised plan is more defensible than a generic activity prescription because it connects the intervention to an observed problem and to the person who must implement it.

Structured psychosocial intervention is better implemented when caregivers are supported. This is not merely because encouragement improves morale. Support protects attention, consistency and willingness to persist through imperfect early attempts. It also gives the caregiver a place to report that the plan is unrealistic or that the behaviour has changed. Behavioural management and caregiver psychoeducation are generally successful, and effects may persist for **months**, but durability still depends on continued fit between the plan, the patient and the caregiver's capacity.

TRAP **Do not write "family counselling" as the psychosocial plan.** A complete answer names the target behaviour, active caregiver role, structured education, skills practice, support, respite and review of feasibility.

The boundary with the BPSD management section matters. This section does not repeat symptom-specific drug choices or their risks. Its contribution is to explain why non-drug treatment is carried through the caregiver system and why unsupported caregivers cannot be expected to deliver complex behavioural plans reliably. When medication is considered elsewhere, caregiver observation remains important for describing the target, detecting change and reporting harm, but pharmacological teaching belongs with BPSD management.

25.6 Management architecture: LOCUS, FOCUS and MODUS

LOCUS asks where care is actually happening and who is available there. A plan for a household, outpatient service or inpatient setting must match its staffing, routines, privacy and capacity for follow-up. The clinician should not export a ward strategy into a home without checking whether the available caregiver can carry it. Movement between settings is a vulnerable point because advice may travel without the practical knowledge of what has worked.

FOCUS changes with the care phase. In an acute difficulty, protect the patient and caregiver while clarifying the immediate source of distress. In the short term, make the most difficult care task manageable and restore a workable routine. In the middle term, consolidate skills, share labour and prevent isolation from deepening. In the long term, sustain the caregiver system, revisit goals and prepare for increasing dependence without pretending that the earlier plan will remain sufficient.

MODUS must remain multicomponent. It includes education, behavioural and psychological work, individual and group support, respite, social and practical assistance, medical review of the patient, and palliative care when the goal changes. These modes are coordinated around a shared formulation. Adding services without deciding what each component is meant to change creates activity but not treatment.

PLANNING AXIS	IMMEDIATE QUESTION	CLINICAL OUTPUT
LOCUS	Where will the plan be used and who can carry it?	A setting-specific, feasible care arrangement
FOCUS: acute	What threatens safety or makes care impossible now?	Containment, clarification and immediate support
FOCUS: short term	Which task or interaction should become manageable?	A defined behavioural and caregiver target
FOCUS: middle term	Which skills and supports need consolidation?	Competency, shared work and respite
FOCUS: long term	How will the plan adapt as dependence grows?	Reassessment, continuity and goal revision
MODUS	Which psychological, social, medical and palliative modes fit?	A coordinated multicomponent plan

Review the original targets. Has the patient's distress changed? Is the caregiver more confident or better supported? Can the household still deliver the intervention, and has respite occurred?

Attendance or information delivery alone does not establish benefit. The outcome lies in the functioning of the care relationship.

25.7 Recognising the transition to end-of-life care

In the final stages of **Alzheimer disease**, patients become bedridden and develop **swallowing impairment**. Pneumonia and other infections commonly become the illnesses through which death occurs. These changes should prompt an explicit shift in clinical attention. The question is no longer only what deficit can be corrected, but what burdens can be relieved and what form of care best preserves comfort and dignity.

The **LOCUS** of end-of-life care must be able to provide comfort and support the family. **FOCUS** moves toward distress, bodily comfort, communication and the emotional needs surrounding dying. **MODUS** becomes palliative and family-centred while retaining ordinary clinical attentiveness. The shift should not be interpreted as withdrawal of care. It is a change in the purpose by which each intervention is judged.

Communication is central because caregivers may experience the transition as failure, abandonment or a demand that they make impossible decisions. The clinician should explain the observed change, acknowledge uncertainty, elicit the patient's known values and help the family understand the goal of comfort. Language should be direct enough to permit preparation and gentle enough to leave room for grief. Avoid presenting a complex decision as though there were a morally correct answer obvious to every family.

> **PITFALL**
>
> Do not continue a restorative style of care by inertia after the clinical goal has become comfort. Equally, do not use the word "palliative" to justify inattentive care. The plan still requires observation, explanation, symptom relief and active family support.

The caregiver's role also changes. Earlier, the emphasis may have been on learning techniques and sustaining routines. Near the end of life, the caregiver may need permission to relinquish performance goals, help in recognising comfort, and support while witnessing decline. Clinical review should ask not only what the family can do, but what should no longer be demanded of them.

25.8 Palliative care, dignity and support for the family

Palliative care in end-stage dementia is helpful because it organises treatment around comfort and supports the family in seeking a peaceful and dignified death. Dignity is expressed through the manner of care: explaining before touch, preserving privacy, responding to distress, avoiding unnecessary disruption and continuing to address the person as a person. Severe cognitive impairment does not remove relational needs or the ethical importance of ordinary respect.

The palliative plan should be coherent across clinicians and caregivers. Each proposed intervention is considered against the current goal, its likely burden and the comfort it may offer. Decisions should not be fragmented into isolated technical choices while the family remains unclear about the overall direction. A shared account of the goal allows different members of the

team to communicate consistently and prevents the caregiver from repeatedly carrying the same explanation between services.

Family support includes anticipatory explanation, space for grief and ambivalence, and recognition of the long period of labour that may precede death. Relief, sadness, guilt, exhaustion and affection can coexist. The clinician need not force these reactions into a uniform narrative. The therapeutic task is to validate the complexity, correct avoidable misunderstanding and make sure that the caregiver is not left alone with practical and emotional responsibility.

After the goal has shifted, caregiver support remains structured. Ask what the family understands, what they fear, who needs to be involved, what care is feasible in the chosen locus, and where additional help is required. Revisit the plan as the patient's condition changes. Consistency does not mean rigidity; it means that adaptations remain visibly tied to comfort, dignity and family support.

The principle is continuous. Caregivers must be noticed, educated, equipped and relieved. Psychosocial treatment works through support, not instructions alone. In advanced disease, the focus moves toward comfort, family support and a dignified death. The caregiver is never merely an informant, and palliative care is not merely absence of curative intent.

> Support the caregiver as a person under strain and as the key treatment partner; when dementia reaches its end stage, reorganise care around comfort, family support and dignity.

26 · Disease-modifying anti-amyloid antibodies: lecanemab, donanemab and the aducanumab debate

Why this topic matters. Anti-amyloid antibodies have changed the Alzheimer treatment question from whether plaques can be removed to whether removing them produces enough clinical benefit to justify infusion, biomarker confirmation, imaging surveillance and haemorrhagic risk. The examination marks therefore lie in interpretation, not in merely listing drug names. A strong answer separates biological action from patient-important benefit, identifies the early symptomatic population in which benefit was studied, explains amyloid-related imaging abnormalities, and uses the regulatory history of aducanumab to show why plaque clearance alone cannot settle efficacy.

These drugs are not conventional symptomatic cognitive enhancers. They intervene in the pathological amyloid pathway and may slow deterioration in selected patients, but they do not imply restoration of lost cognition or a cure. Their use consequently begins before a prescription is written: the clinical syndrome, stage and amyloid evidence must fit the evidence base, while the patient and family must understand the scale and uncertainty of benefit as well as the monitoring burden.

Target	Entry	Outcome	Constraint
AMYLOID	BIOMARKER PROOF	SLOWER DECLINE	ARIA

26.1 Therapeutic logic: from amyloid biology to immunotherapy

Anti-amyloid therapy rests on a pathological sequence in which secretase cleavage of amyloid precursor protein yields beta-amyloid, pathological beta-amyloid accumulates and plaques contain this material; therapeutic development has consequently pursued secretase inhibition, inhibition of beta-amyloid aggregation and beta-amyloid immunotherapy. The antibodies discussed here belong to the immunotherapy strategy. This classification is useful because it prevents the common error of treating all "anti-amyloid" agents as pharmacologically interchangeable.

An antibody can engage a particular physical form of beta-amyloid and reduce plaque burden. That is a biological effect. Disease modification, however, is a clinical claim: the treated group should deteriorate less than an appropriate comparator on outcomes that reflect cognition or function. These are related claims, but they are not synonyms. Plaque removal supports target engagement; it does not by itself tell the clinician how much longer a person will preserve independence, whether an individual patient will perceive improvement, or whether risk outweighs benefit.

RULE **Separate target engagement from clinical benefit.** Amyloid reduction answers "did the drug affect its target?"; comparative cognitive and functional outcomes answer "did that biological effect matter clinically?"

The same distinction shapes consent. "Slowing decline" means a difference in trajectories between groups while disease continues to progress. It should not be translated into "memory improves" or "dementia is reversed." A modest average separation can still matter to some patients, but its value depends on baseline stage, goals, treatment burden, imaging risk and the uncertainty of applying a trial mean to an individual patient.

26.2 Lecanemab: target, evidence and interpretation

Lecanemab lowers brain beta-amyloid plaque burden by binding soluble beta-amyloid protofibrils, with variable binding to other beta-amyloid forms. The molecular target matters because the drug is not simply a nonspecific plaque solvent. Its rationale is directed engagement of an amyloid species within a pathway that culminates in plaque accumulation.

The pivotal evidence enrolled people with mild cognitive impairment or early Alzheimer disease whose amyloid pathology had been confirmed by positron-emission tomography or cerebrospinal fluid. In that setting, **the trial included 1795 participants and compared lecanemab 10 mg/kg by intravenous infusion every 2 weeks with matched placebo**. This sentence carries the clinically important scope: the dose belongs to early symptomatic Alzheimer disease with amyloid confirmation, not to unselected cognitive complaints or advanced dementia.

On the **Clinical Dementia Rating Sum of Boxes**, **lecanemab reduced decline by 27% relative to placebo over 18 months, an absolute difference of 0.45 points, with changes of 1.21 and 1.66 respectively**. Key secondary outcomes also favoured treatment. The relative figure sounds larger than the absolute separation because they express different things. The relative result describes proportional slowing against the comparator trajectory; the absolute

result shows how far apart the groups actually were on the scale. An exam answer should state both concepts and should never convert slowing into improvement.

The appropriate reading is therefore neither dismissal nor triumphalism. The trial demonstrated amyloid reduction accompanied by statistically supported slowing of measured decline in the studied population. It did not establish that every treated patient notices benefit, and it does not justify extending the claim beyond the stage and biomarker conditions studied. Selection is part of the treatment, not an administrative prelude to it.

GOOD TO REMEMBER

When explaining lecanemab, separate the message into clear statements: amyloid is confirmed; the drug slows average decline rather than improving established deficits; imaging risk and infusion burden remain part of the treatment decision.

26.3 Donanemab: plaque clearance and clinical slowing

Donanemab is a high-potency anti-amyloid treatment given by intravenous infusion whose pivotal study concerned early symptomatic Alzheimer disease, including mild cognitive impairment and mild dementia, with amyloid pathology and tau-pathology strata demonstrated on positron-emission tomography. This reinforces the class lesson: biological confirmation and early clinical stage define the evidentiary population.

Compared with placebo, **donanemab slowed cognitive and functional decline by approximately 35% over 18 months in the primary target population, while 52% of treated participants became amyloid-PET negative by 12 months**. The pairing of these results is instructive. Conversion to an amyloid-negative scan is a striking marker of target effect; slowing on cognitive and functional outcomes is the clinically relevant companion result. The scan outcome cannot be substituted for the clinical outcome, even when both move in the expected direction.

Donanemab also illustrates why a treatment programme is more than the act of infusion. The clinician must establish that the syndrome and stage correspond to those studied, demonstrate amyloid pathology, interpret tau information where it forms part of the treatment pathway, obtain suitable baseline imaging, discuss realistic benefit and arrange continuing safety surveillance. The allowed evidence supports an intravenous schedule but does not supply a numeric maintenance amount; that absence should remain explicit rather than being filled from memory.

Comparing donanemab with lecanemab by placing headline percentages side by side is tempting but methodologically weak. Different trial populations, outcome frameworks and biomarker selection can change apparent effect sizes. The safe conclusion is that each antibody showed slowing against its own placebo group in an early, biomarker-confirmed population. A direct claim that either is clinically superior requires comparative evidence not contained in these trial summaries.

26.4 ARIA: the class-defining imaging hazard

Amyloid-related imaging abnormalities, abbreviated ARIA, refer to brain swelling or bleeding seen in association with these treatments. The useful clinical distinction is between oedematous change and haemorrhagic change. ARIA may be discovered on surveillance imaging even without symptoms, but it can also become clinically important. The possibility of an asymptomatic lesion is exactly why symptom-triggered scanning alone is an inadequate conceptual model.

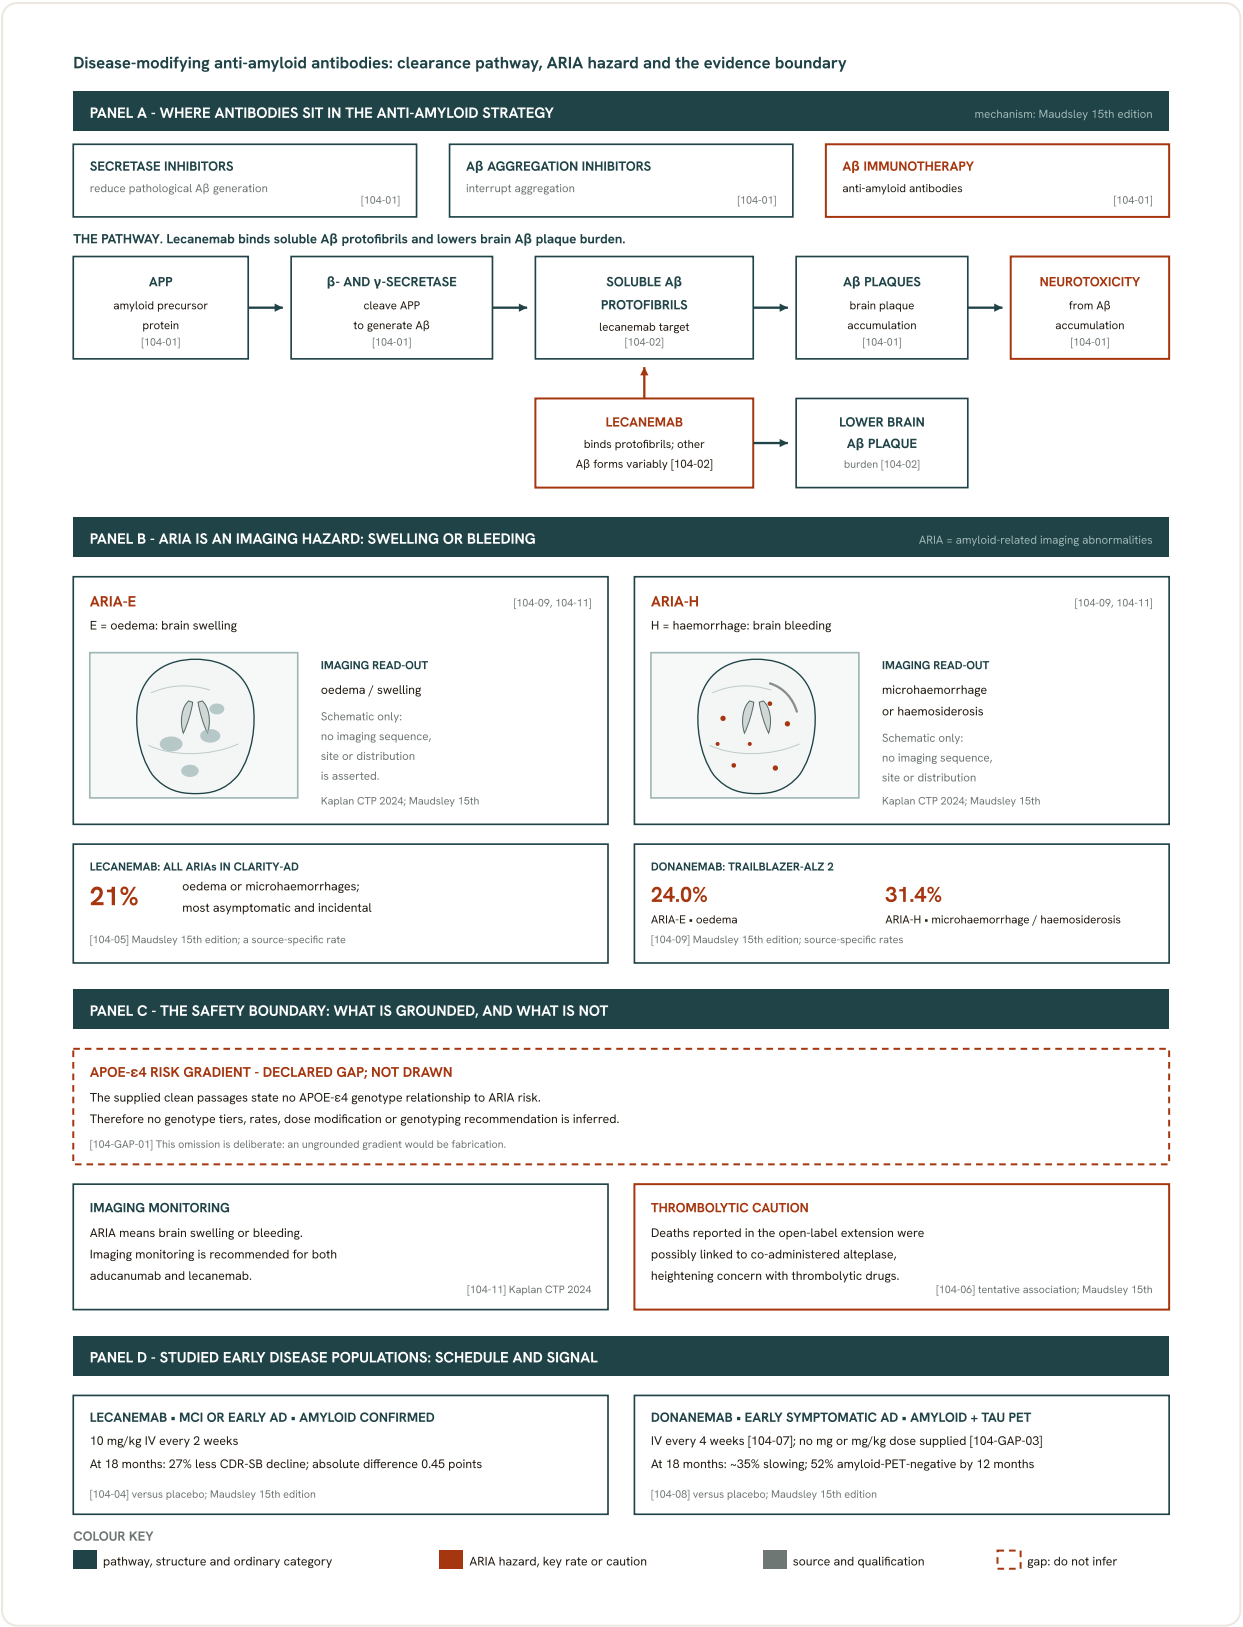


Fig 32.11 Anti-amyloid antibodies and amyloid-related imaging abnormalities (ARIA). Lecanemab binds soluble Aβ protofibrils and lowers plaque burden; ARIA-E denotes oedema and ARIA-H denotes haemorrhagic findings. The rates shown remain attached to their own trial and source rows. The requested APOE-ε4 risk gradient is not drawn because the supplied evidence contains no genotype-risk relationship. (Maudsley Prescribing Guidelines in Psychiatry, 15th ed.; Kaplan and Sadock's Comprehensive Textbook of Psychiatry, 2024.)

The accompanying figure links amyloid clearance with the distinct imaging patterns. ARIA-E denotes the oedematous pattern, while ARIA-H denotes microhaemorrhage or haemosiderosis. In

the donanemab study, **ARIA-E occurred in 24.0% and ARIA-H in 31.4% of treated participants**. These categories are not interchangeable, and their percentages should not be added as if they represented mutually exclusive patients.

With lecanemab, ARIA affected a substantial minority of the treated group; **asymptomatic radiological detection** accounted for most reported cases. This is a practical warning against equating "no complaint" with "no adverse effect." Imaging surveillance is recommended for lecanemab and aducanumab because ARIA is fundamentally an imaging-defined safety problem. Exact scan timing, thresholds for interruption and criteria for resumption must come from the applicable current protocol rather than from an unsourced generic schedule.

A separate safety signal concerns **thrombolytic co-administration**: deaths reported during open-label extension treatment were considered possibly linked to co-administration of alteplase. The wording must remain cautious: this is a possible association, not proof that alteplase caused every event. Nevertheless, acute-care teams need to know that a patient is receiving anti-amyloid therapy when evaluating neurological deterioration or a reperfusion decision.

PITFALL

Do not reassure solely because the patient is asymptomatic. ARIA can be an incidental imaging finding, so safe treatment depends on planned surveillance and a clear pathway for responding to new neurological symptoms or radiological change.

MNEMONIC

Plaque proof, Benefit modest, Radiology matters

Plaque proof defines entry into the evidence base; **Benefit modest** prevents "slowing" from becoming "reversal"; **Radiology matters** keeps ARIA surveillance inside the treatment plan.

26.5 Comparative frame: what can and cannot be concluded

QUESTION	LECANEMAB	DONANEMAB	ADUCANUMAB
Amyloid action	Binds soluble protofibrils and lowers plaque burden	Potent anti-amyloid action with substantial plaque clearance	Preferentially binds aggregated beta-amyloid and reduces plaques
Evidence population	Early symptomatic disease with amyloid confirmation	Early symptomatic disease with amyloid and tau imaging characterisation	Mild cognitive impairment or mild dementia due to Alzheimer disease
Clinical evidence	Slower decline than placebo on cognitive and secondary outcomes	Slower cognitive and functional decline than placebo	Conflicting clinical-decline results underlay the regulatory dispute
Biomarker evidence	Reduced brain amyloid	Many treated participants reached amyloid-negative imaging status	Plaque reduction formed the surrogate basis of approval
Central safety issue	ARIA, often detected without symptoms	Both oedematous and haemorrhagic ARIA patterns	ARIA drove concern about the risk-benefit balance
Core interpretation	Target engagement plus demonstrated slowing in the studied population	Marked plaque effect plus demonstrated slowing in the studied population	Target engagement did not resolve uncertainty about clinical benefit

The table should be read across rows, not as a league table. These agents engage amyloid, yet their evidentiary stories differ. Lecanemab and donanemab are taught through concordance between biomarker change and slower decline against placebo. Aducanumab is taught through discordance: convincing plaque reduction existed alongside disputed clinical evidence. That contrast is the conceptual centre of the debate.

Several conclusions are unsafe. Plaque clearance does not establish cure. A larger relative percentage from a trial does not prove superiority over another drug. Absence of symptoms does not exclude ARIA. Approval does not mean that uncertainty has disappeared. Conversely, controversy surrounding an antibody does not invalidate the entire therapeutic target; each agent requires appraisal of its own target, population, clinical endpoint and safety evidence.

26.6 Aducanumab: mechanism and the approval controversy

Aducanumab targets beta-amyloid and preferentially binds aggregated beta-amyloid, reducing its accumulation and plaque number. After titration it is administered as **maintenance intravenous therapy**. Mechanistically, therefore, it has an intelligible disease-modifying rationale and clear evidence of biological target engagement. The controversy was not about whether amyloid plaques could fall. It was about whether the available clinical evidence justified inferring meaningful benefit and accepting the associated harms.

Development was initially stopped after futility analysis in parallel late-stage trials. The manufacturer's later **reanalysis** included participants who had continued beyond the futility cut-

off, reported a significant finding in part of the programme with supporting evidence elsewhere, and preceded an application for United States approval; conditional accelerated approval followed, while approval was not granted in Europe. This sequence explains why the dispute cannot be reduced to “the trial was positive” or “the trial was negative.” The interpretation changed after additional data and post hoc analysis, while consistency across trials remained a central concern.

The United States decision used reduction of beta-amyloid plaque as a **surrogate endpoint** because clinical-decline findings conflicted. Accelerated approval therefore carried an obligation to verify clinical benefit after approval. It was historically important as a disease-modifying approval, but historical importance and evidentiary certainty are different judgments.

Experts questioned both the extent of clinical benefit and the **risk-benefit profile**, particularly because ARIA may include vasogenic oedema, sulcal effusion or microhaemorrhage; the drug was directed to mild cognitive impairment and mild dementia due to Alzheimer disease, not to dementia of any cause or severity. The scientifically fair position is neither that plaque reduction is worthless nor that it is sufficient. It is useful evidence about mechanism, but clinical benefit must be demonstrated on clinical outcomes.

TRAP **Do not turn an internally variable approval date into false precision.** Parallel textbook passages describe lecanemab accelerated approval as **January 2023** and as early in the same year. Preserve that wording difference, distinguish accelerated from later traditional approval, and focus the answer on the confirmatory demonstration of clinical benefit.

26.7 Clinical deployment: LOCUS, FOCUS and MODUS

LOCUS. Initiation belongs in a specialist outpatient service able to confirm the syndrome and amyloid pathology, deliver intravenous therapy, review brain imaging and coordinate urgent assessment. Hospital or emergency care becomes the relevant locus when new neurological symptoms raise concern for oedema or bleeding. The service must communicate treatment status across memory, radiology, neurology and acute-care teams because an infusion decision made in a particular setting can alter risk assessment elsewhere.

FOCUS. Before treatment, the focus is eligibility: early symptomatic Alzheimer disease, evidence of amyloid pathology, a baseline clinical and imaging picture suitable for later comparison, and goals realistic enough for informed consent. During treatment, the focus shifts to preserving function by slowing decline, maintaining adherence to infusions and surveillance, and detecting imaging or neurological adverse effects. Over longer follow-up, the focus is whether ongoing burden and risk remain proportionate to observed course and patient priorities.

MODUS. The mode is not drug-only care. It combines intravenous antibody treatment, biomarker-supported selection, serial clinical review, imaging surveillance, adverse-event pathways, education and shared decision-making. Supportive dementia care continues; disease modification does not replace attention to function, comorbidity, caregiver needs, safety or advance planning. Current jurisdiction-specific product information and clinical guidance must

supply exact imaging schedules, interruption rules and eligibility exclusions because these operational details are not interchangeable between antibodies.

- **Before starting:** confirm that the clinical stage and amyloid evidence match the studied population; document baseline cognition, function and imaging; discuss slowing rather than recovery.
- **During treatment:** preserve infusion and imaging continuity; ask about new neurological symptoms; ensure radiological findings reach the prescribing team promptly.
- **Across settings:** make anti-amyloid treatment visible to acute-care clinicians, especially when haemorrhage or oedema enters the differential; the lecanemab safety signal with a thrombolytic is a specific caution worth extending prudently to the class, given the shared haemorrhagic risk of ARIA.

Guidelines-first practice here means using the current local label and specialist protocol rather than memorising an attractive universal schedule. The supplied evidence establishes that imaging monitoring is recommended, but it does not provide the complete timing or stop-and-restart algorithm. A safe textbook answer states that boundary explicitly.

26.8 Consent, viva framing and final synthesis

Consent should be organised around decisions rather than data dumping. The patient and family need to understand the treatment aim, the biomarker requirement, the expected direction of benefit, infusion and imaging burden, ARIA, uncertainty about individual response and the possibility that treatment may be paused or stopped. Cognitive impairment makes supported decision-making especially important, but support should clarify the person's values rather than replace them.

A concise viva answer can follow a linked sequence:

1. Define the class as beta-amyloid immunotherapy, then state that lecanemab and donanemab reduced amyloid and slowed decline against placebo in selected early symptomatic, biomarker-confirmed populations.
2. Explain ARIA as oedematous or haemorrhagic imaging abnormalities; because they occur with each of these antibodies, at their highest rates with donanemab, imaging surveillance is integral to treatment, with the monitoring recommendation stated explicitly for aducanumab and lecanemab.
3. Use aducanumab to distinguish surrogate plaque reduction from verified clinical benefit and to explain accelerated-approval controversy.

This structure also prevents a fragmented answer. Mechanism leads to selection, selection leads to efficacy, efficacy must be balanced against ARIA, and the approval debate tests whether the candidate understands the hierarchy of evidence. If time is short, preserve that chain rather than reciting isolated brand facts. It demonstrates both pharmacology and clinical judgment.

The mature conclusion is balanced. Anti-amyloid antibodies are a genuine advance because biological target engagement has been accompanied, for some agents, by slower cognitive or functional decline in the populations studied. Their benefits remain modest at group level, their administration is resource-intensive, and imaging harm is central rather than incidental. The correct candidate is therefore not “anyone with dementia” but a person whose phenotype, stage, biomarker status, safety profile, care setting and preferences fit the evidence and the applicable protocol.

> Anti-amyloid antibodies may slow decline in selected early, amyloid-confirmed Alzheimer disease, but plaque clearance is not cure and ARIA surveillance is part of the treatment itself.

27 · Down syndrome and early-onset Alzheimer's disease: chromosome 21, amyloid and assessment of decline in intellectual disability

The marks lie in connecting chromosome, gene dosage, tissue change and acquired dementia without collapsing them into one claim. **Down syndrome** provides a clinically important bridge between a developmental chromosomal disorder and an Alzheimer-like neurodegenerative syndrome. The bridge is not merely an association. The cited account links the chromosomal substrate to extra copies of a gene involved in amyloid biology, and links the later clinical syndrome to plaques, tangles and amyloid accumulation. A strong answer preserves every step in that chain and also preserves the uncertainty around the proposed mechanism.

The diagnostic problem is equally important. Intellectual disability is a pre-existing developmental condition, whereas dementia denotes an acquired clinical disorder. The chromosomal diagnosis, the proposed molecular mechanism, the presence of characteristic tissue changes and the clinical diagnosis answer different questions. None should be used as a shortcut for another. Intellectual disability and Down syndrome as a developmental syndrome are taught in the Intellectual disability and autism spectrum disorder notes; this section remains focused on the later dementia relationship.

Chromosome 21	3 APP COPIES	BY 50 YEARS	6D85.9
CHROMOSOMAL SUBSTRATE	PROPOSED GENE-DOSAGE LINK	ALMOST UNIFORM COMPLICATION	ICD-11 CATEGORY

27.1 Chromosomal basis: state the cause precisely

The central genetic fact is involvement of chromosome 21. In the cited clinical-neurology account, **90%** of cases arise from **nondisjunction during gametogenesis**, while a **chromosomal translocation** is responsible in **5% or less**. These figures should be quoted with their mechanisms attached. A percentage without the corresponding chromosomal process is inert recall; a process named without its stated scope may falsely imply that it explains every case.

Nondisjunction and translocation are therefore alternative causal routes within the supplied account, not interchangeable words for the same event. The safe formulation is that the developmental syndrome is caused by a chromosomal abnormality involving chromosome 21,

with nondisjunction during gamete formation accounting for the great majority and translocation for a much smaller proportion. The supplied figures do not form an exhaustive genetic catalogue. Their total should not be manipulated to invent an unnamed residual category or an unsupported percentage.

This precision matters later. The amyloid theory depends on increased dosage of a particular gene situated on the same chromosome. The chromosomal cause and the later molecular hypothesis are connected, but they remain distinct propositions. The causal statement concerns why the developmental syndrome occurs. The gene-dosage statement proposes why people with that syndrome tend to develop an Alzheimer-like dementia. Keeping those propositions separate prevents circular reasoning.

RULE **Move from chromosome to gene dosage to amyloid to the clinical syndrome.** Each link answers a different question, and no link by itself proves the whole chain in an individual patient.

27.2 Parental age: association is not individual causation

The incidence of the developmental syndrome correlates with increasing **maternal age**, with the association especially noted **after 40 years**. Increasing paternal age is also associated with greater incidence, but the relation is less marked than the maternal-age association. The direction of comparison must be retained: paternal age is not stated to have an equal effect, and maternal age is not presented as the only population correlate.

These are incidence associations. They describe how occurrence varies across groups and should not be rewritten as a deterministic rule for a particular pregnancy or family. Maternal age is neither the chromosomal mechanism itself nor a retrospective assignment of responsibility. The mechanism in the supplied account is nondisjunction of chromosome 21 during gametogenesis in most cases, or translocation in a smaller group. Parental age and chromosomal mechanism belong at different explanatory levels.

The same distinction protects later dementia reasoning. A factor associated with occurrence of Down syndrome should not automatically be presented as a factor governing the timing or certainty of the later dementia. The allowed evidence supports an age association for incidence of the developmental condition and a separate age-linked statement about the Alzheimer-like complication. It does not establish that parental age predicts the later cognitive course, neuropathological burden or clinical severity. That inference must not be smuggled into an otherwise correct genetic answer.

PITFALL

Do not use a population association to assign blame to parents or to predict the later dementia trajectory. State the maternal-age association, acknowledge the weaker paternal-age association, and then return to the chromosomal mechanism actually supplied.

27.3 The amyloid gene-dosage hypothesis

The mechanistic centre of the topic is the **amyloid precursor protein gene**. The cited source frames the explanation as one theory: people with Down syndrome have **three copies** of the gene, which is situated on chromosome 21, and this extra gene dosage is proposed to cause **amyloid overproduction**. The wording “one theory” is not decorative caution. It is part of the factual claim and must survive into the answer.

The logic is attractive because it creates a continuous biological narrative. The chromosomal disorder changes gene copy number; the relevant gene lies on the affected chromosome; greater dosage is proposed to increase amyloid production; and the later dementia is accompanied by amyloid accumulation. This is a mechanism-level explanation, not merely a juxtaposition of facts. It explains why the association has particular importance for understanding the wider disease biology.

Nevertheless, coherence is not the same as complete proof. The allowed evidence does not say that gene dosage is the sole pathway, that amyloid alone is sufficient for clinical dementia, or that the clinical timing can be calculated from copy number. It also supplies no mutation catalogue, laboratory threshold or individual predictive test. A disciplined answer therefore says “one theory holds” or “the proposed mechanism is” rather than presenting the chain as an exclusive settled cause.

The direction of reasoning also matters. The theory begins with the known chromosomal state and proposes a downstream consequence for amyloid production. It should not be reversed into the claim that finding amyloid proves Down syndrome, or that any Alzheimer-like presentation must arise through the same chromosomal route. The model explains a particular association in a defined population. It is not a universal diagnostic rule for dementia and it does not erase the need to establish the clinical syndrome.

Gene dosage, pathology and phenotype are therefore related but non-equivalent levels of explanation. Gene dosage is the proposed upstream driver. Amyloid overproduction is the proposed molecular consequence. Plaques, tangles and amyloid accumulation are the associated tissue findings. Alzheimer-like dementia is the clinical phenotype. An examination answer becomes more accurate, not less confident, when it names these levels and keeps the source's uncertainty attached to the mechanistic step.

EXPLANATORY LEVEL	WHAT THE SUPPLIED EVIDENCE SUPPORTS	WHAT MUST NOT BE INFERRED
Chromosomal cause	Down syndrome usually follows nondisjunction involving chromosome 21; translocation accounts for a smaller proportion	That the listed mechanisms exhaust every possible case
Gene dosage	The amyloid precursor protein gene is on chromosome 21, and affected individuals have an extra copy	That copy number alone diagnoses dementia
Proposed molecular effect	Extra gene dosage is theorised to increase amyloid production	That the theory is the only mechanism or predicts an individual's course
Tissue phenotype	The later syndrome is associated with plaques, tangles and amyloid accumulation	That a qualitative pathological association supplies a clinical prevalence estimate
Clinical phenotype	An Alzheimer-like dementia becomes an almost uniform complication by the age stated in the source	That every person at a younger age already has clinical dementia
Classification	ICD-11 lists dementia due to Down syndrome as a discrete category	That a category label substitutes for clinical assessment

MNEMONIC

CAPS

Chromosome 21; **A**myloid precursor protein gene dosage; **P**laques and tangles with amyloid accumulation; **S**yndrome of Alzheimer-like dementia. CAPS preserves the causal sequence while leaving room for the source's “one theory” qualification.

27.4 From amyloid accumulation to Alzheimer-like dementia

The clinical-neuropathological link is unusually direct in the supplied account. Down syndrome leads to an **Alzheimer-like dementia** associated with **cerebral plaques**, **neurofibrillary tangles** and accumulation of amyloid. The combination matters. Amyloid is not presented as an isolated laboratory observation; it sits beside characteristic structural lesions and a clinical dementia phenotype.

The temporal statement is also strong: the dementia is described as an almost uniform complication **by age 50 years**. “Almost uniform” should not be silently converted into a percentage, and the age statement should not be converted into an earlier neuropathological threshold. The source gives a qualitative description of near-universality at the specified age and a qualitative account of the associated pathology. It does not provide prevalence by decade or a numerical separation between tissue change and clinical dementia.

This distinction is central to examination answers. A candidate may remember a widely repeated age for pathological change or a prevalence estimate for clinical dementia, but those figures are not supplied here. The defensible answer is narrower and stronger: the later dementia is almost uniform by the stated age and shows plaques, tangles and amyloid accumulation. Nothing is gained by adding an unverified earlier age or an uncited percentage.

The pathological terms should also remain together. Plaques and tangles describe structural lesions, while amyloid accumulation supplies the molecular emphasis that connects back to the gene-dosage theory. Selecting amyloid alone makes the account sound more complete mechanistically than the source permits; listing pathology without the dementia phenotype loses the clinical point. The supported formulation is a linked clinical-pathological syndrome, not an isolated protein abnormality.

The phrase “Alzheimer-like” also deserves care. It communicates close clinical and pathological resemblance while remaining the wording of the cited source. It should not be used to re-teach the entire neuropathology, staging, biomarker profile or treatment of the wider disease. Those constructs are developed in their dedicated sections. Here, the educational point is the specific bridge from the chromosomal disorder to amyloid-related dementia.

TRAP **Do not turn a strong association into an invented prevalence table.** “Almost uniform by the stated age” does not license a clinical percentage for each decade, and qualitative neuropathology does not license an earlier numerical threshold for plaques or amyloid.

27.5 The assessment problem in pre-existing intellectual disability

Assessment is the clinically hardest part, but the supplied evidence does not provide an operational method. The conceptual distinction is clear: the developmental diagnosis and the acquired dementia diagnosis are not synonyms. The available claims establish the genetic substrate, a proposed amyloid mechanism, the later clinical-pathological syndrome and its diagnostic-manual category. They do not establish how decline should be measured against pre-existing intellectual disability.

Accordingly, this fragment does not name a Down-syndrome-specific dementia instrument, reproduce an informant questionnaire, print a threshold, prescribe an assessment interval or provide a rule for interpreting serial scores. It also does not claim that a generic cognitive screen has been validated for this purpose. General cognitive-screen administration and scoring belong in the Psychometry, Assessment and Descriptive Psychopathology notes, but that cross-reference does not fill the missing Down-syndrome-specific evidence.

What can safely be said is a limit on inference. A chromosomal diagnosis demonstrates the developmental condition; it does not by itself demonstrate acquired dementia. The gene-dosage theory supplies biological plausibility; it does not establish current functional decline in an individual. The strong age-related association changes clinical concern; it does not replace assessment. Plaques, tangles and amyloid accumulation explain the pathological resemblance; they do not provide an allowed bedside algorithm in this evidence set.

GOOD TO REMEMBER

When the evidence bundle does not supply an assessment instrument or threshold, do not manufacture one from memory. State the developmental condition and the suspected acquired dementia separately, then identify the need for an appropriate decline assessment without pretending that an unsupported tool or cut-off has been established here.

27.6 Diagnostic formulation and ICD-11 classification

ICD-11 recognises **dementia due to Down syndrome** as a discrete category within the dementia block, coded 6D85.9. The classification has useful implications. It acknowledges a specific etiological relationship rather than forcing the presentation into an unspecified dementia label, and it keeps dementia conceptually distinct from the developmental syndrome itself.

The code should be attached to the complete construct, not memorised as a floating number. “Due to Down syndrome” identifies the etiological attribution, while placement within the dementia block identifies the acquired clinical disorder being classified. The developmental syndrome is not being recoded as dementia. Conversely, the availability of a dedicated category does not mean that every adult with Down syndrome automatically meets the clinical requirements for dementia at every age.

The allowed claim establishes category and code, not the full diagnostic criteria. It does not provide required symptom domains, duration, severity thresholds, exclusion rules or a crosswalk to another diagnostic manual. These should not be reconstructed from memory and placed under the code. In a written answer, use the category to complete the formulation after explaining the genetic and pathological relationship, not as a substitute for that explanation.

A clean etiological formulation therefore has layers. Name the established developmental chromosomal disorder. State that the later syndrome is Alzheimer-like and associated with plaques, tangles and amyloid accumulation. Present extra copies of the amyloid precursor protein gene on chromosome 21 as one proposed mechanism for amyloid overproduction. Then, when the clinical diagnosis has actually been established, use the discrete ICD-11 category. The sequence is more informative than code recall alone.

27.7 Building the final examination answer

A high-quality answer begins with chromosome rather than with a vague statement that dementia is “common”. It gives the supplied chromosomal mechanisms with their properly scoped proportions, then introduces the parental-age association without turning it into personal blame or a dementia predictor. It next explains the amyloid precursor protein gene-dosage theory, explicitly retaining the source's theoretical framing.

The answer then moves to phenotype. It states that Alzheimer-like dementia is an almost uniform complication by the specified age and names the associated plaques, tangles and amyloid accumulation. It distinguishes this supported statement from the unsupported claim that all affected adults have clinical dementia at an earlier age. Finally, it names the dedicated diagnostic-

manual category and acknowledges that the supplied evidence does not furnish an operational protocol for measuring decline in pre-existing intellectual disability.

- **Genetic cause:** name the chromosome 21 abnormality and keep nondisjunction separate from translocation.
- **Epidemiological association:** state the stronger maternal-age relation and the weaker paternal-age relation without extending either to dementia prognosis.
- **Mechanism:** present extra amyloid precursor protein gene dosage as one theory of amyloid overproduction.
- **Phenotype:** link the Alzheimer-like dementia to plaques, tangles and amyloid accumulation, preserving the source's age wording.
- **Diagnostic boundary:** separate developmental syndrome, biological plausibility, tissue phenotype and acquired dementia classification.

The common errors all come from collapsing levels. A candidate may treat parental age as the mechanism, treat the chromosomal diagnosis as proof of dementia, treat gene dosage as a bedside diagnostic test, treat pathology as identical to clinical prevalence, or treat a classification code as complete criteria. The cure is the same each time: state what kind of claim is being made and stop at the limit of the evidence.

That restraint is not thin knowledge. It is the core clinical method in a high-fabrication-risk topic. The causal chain is memorable precisely because it is coherent, but coherence can tempt the writer to fill missing links with familiar figures or named instruments. The mature answer explains the supported links in depth and declares the missing assessment evidence rather than disguising memory as authority.

> Down syndrome links chromosome 21 to an extra amyloid precursor protein gene dose and proposed amyloid overproduction, but dementia remains a distinct acquired clinical diagnosis.

Consolidate and apply

28 · Worked vignettes

A vignette rewards the candidate who can rank evidence, not the candidate who can repeat the stem. Cognitive-disorder cases are deliberately crowded with memory complaints, behavioural observations, medical comorbidity, scan language and family anxiety. The task is to identify which facts establish an acquired cognitive syndrome, which facts define its phenotype and tempo, which facts merely modify performance, and which fact changes what must be done now. A polished answer therefore moves from syndrome to cause to action while preserving uncertainty.

These cases consolidate the chapter without replacing its disease-specific sections. Each vignette is followed by the reasoning that should be audible in a written answer or viva: the diagnostic pivot, the nearest alternative, the missing evidence and the immediate care consequence. Exact criteria, biomarker cut-offs, drug doses and monitoring schedules must be retrieved from their dedicated sections rather than reconstructed from memory.

28.1 Decode the stem before choosing a label

Read a cognitive vignette in passes. Begin by asking whether the stem demonstrates decline from a previous level or only reports a complaint. Identify the dominant cognitive, behavioural or neurological pattern. Establish tempo from concrete changes rather than adjectives such as "recent" or "gradual." Decide whether everyday independence has changed because of cognition. Close the read by identifying a safety issue, active contributor or discordant fact that could alter urgency.

Tempo	Pattern	Function	Safety
SETS THE ROUTE	RANKS CAUSES	SETS THE LEVEL	SETS THE ACTION

The **diagnostic pivot** is the feature, or coherent cluster of features, that most changes the ranking of plausible explanations. It is not always the most dramatic symptom. Repeated forgetting may be less discriminating than the way information is lost, the sequence of behavioural change, genuine fluctuation, focal neurological evidence or a mismatch between reported disability and observed function. State the pivot explicitly. This shows that the answer is reasoned rather than guessed.

A useful opening sentence contains the syndrome, phenotype, tempo and functional effect. The etiological formulation follows and should be qualified by the strength of evidence. Then name the strongest alternative and the observation that would separate it. If the question asks for the next step, do not bury that step beneath a complete differential. Answer the decision being tested.

RULE **Every vignette answer should contain a pivot, a competitor and a consequence.** Say what finding leads, what plausible explanation still competes, and what the formulation changes in assessment, urgency or care.

28.2 Vignette: progressive memory-led decline

A retired teacher is brought because conversations are repeated, recent events are lost and familiar household responsibilities have gradually passed to the spouse. Remote autobiographical material remains readily available. The patient gives plausible explanations for mistakes and regards the family's concern as excessive. Examination shows poor acquisition and retention of new information, while social manner and basic language remain relatively preserved. The history contains neither abrupt change nor prominent fluctuation or a focal neurological event.

The safest answer begins with a **progressive amnestic neurocognitive syndrome**, not an etiological name. The stem supports acquired decline through collateral examples, a memory-led phenotype through the error pattern, and clinically meaningful functional effect through transfer of previously managed responsibilities. Reduced awareness strengthens the need for collateral

evidence but is not itself diagnostic. The absence of an abrupt or fluctuating course makes some alternatives less persuasive, yet it does not establish the pathology.

The next reasoning step is to ask whether memory failure reflects inefficient attention or retrieval, impaired storage, language difficulty, mood-related interference, sensory limitation, medication burden or a broader degenerative process. Observe learning across attempts, delayed retention, recognition, cueing and intrusion errors as a pattern rather than reducing cognition to a total score. General administration and scoring of cognitive instruments belong in the Psychometry, Assessment and Descriptive Psychopathology notes.

A strong answer also asks what has been taken over and why. A spouse may compensate so effectively that outward independence looks preserved. Conversely, retirement or traditional division of household roles may mean the patient never performed the task under discussion. Function must be compared with the person's own baseline and linked to cognition. Investigations should then test the leading formulation, identify contributors and determine whether disease-specific confirmation would change counselling or treatment.

GOOD TO REMEMBER

When a stem says “memory loss,” translate it into an error mechanism. Failure to register, rapid loss after learning, poor retrieval with useful cueing, language-based misunderstanding and inconsistent engagement are not interchangeable. Describe the observed pattern before assigning meaning to it.

28.3 Vignette: fluctuation, perception and motor change

An older adult has periods of strikingly clear conversation alternating with episodes of marked cognitive inefficiency that are not explained by sleep or an acute medical state. The family reports recurrent, well-formed visual experiences and increasingly slow, rigid movement. Memory complaints are present, but visuospatial and attentional failures dominate daily mishaps. A previous sedating prescription produced a disproportionate deterioration in mobility and alertness.

The **cortical synucleinopathy phenotype** is raised by the convergence of cognitive fluctuation, recurrent formed visual phenomena, visuospatial-attentional difficulty and spontaneous extrapyramidal change. An isolated feature should not carry the diagnosis. Fluctuation must be distinguished from ordinary good and bad periods, fatigue, medication timing, sleep disturbance, seizures and delirium. Visual experiences require phenomenological description and assessment of vision, context, distress and insight. Motor findings require direct examination rather than reliance on a family adjective such as “slowness.”

The medication history is a safety pivot. It should prompt review of the drug, the target symptom, temporal relation, concurrent illness and recovery after withdrawal. It does not justify transferring a general antipsychotic algorithm into dementia care. General antipsychotic pharmacology belongs in the Schizophrenia: management, antipsychotics and adverse effects

notes; dementia-specific behavioural prescribing and sensitivity are developed in the relevant chapter sections.

The nearest competitors include an amnestic degenerative disorder with later perceptual symptoms, a primary visual or psychiatric explanation, medication effects, sleep-related phenomena and a superimposed acute confusional state. The distinction rests on chronology, clustering, examination and longitudinal corroboration. Movement-disorder neuropsychiatry without dementia belongs in the Neuropsychiatry of systemic and neurological disease notes. In this vignette, the examination question is the dementia phenotype and its prescribing consequence.

TRAP **Do not let a vivid hallucination become the diagnosis.** The discriminating evidence is the coherent longitudinal cluster, and the clinically important consequence is caution with sedating or dopamine-blocking treatment while causes of deterioration are reassessed.

28.4 Vignettes with executive, behavioural or language-led presentations

A patient develops slowing, impaired planning and loss of organisational ability after recurrent cerebrovascular events, with focal neurological signs and an uneven course. A different patient retains orientation and routine memory but becomes tactless, impulsive, emotionally indifferent and rigid in habits, with progressive failure of social judgement. Another speaks fluently yet increasingly loses word meaning and conceptual knowledge, while practical behaviour initially remains comparatively organised.

These stems test whether the candidate can resist the reflex that all dementias are primarily amnestic. The vascular stem suggests a **vascular executive syndrome** when cognitive change, neurological evidence and vascular chronology cohere. A vascular lesion on imaging is not enough by itself; common abnormalities may be contributory, incidental or insufficient to explain the phenotype. The broader pre-dementia and post-stroke cognitive spectrum is taught in the Stroke, TBI, delirium and focal brain syndromes notes.

The behavioural stem suggests a **frontotemporal syndrome** because progressive change in social cognition, inhibition, empathy and behavioural regulation leads the presentation. A primary psychiatric disorder remains a serious competitor. The answer should therefore seek longitudinal loss from baseline, consistency across settings, stereotyped or compulsive patterns, dietary change, language or motor evolution, neurological signs and imaging concordance. A psychiatric history does not immunise a patient against neurodegeneration, and a late behavioural change does not automatically prove it.

The language-led stem requires the candidate to identify which part of language is failing. Fluency, grammar, motor speech, naming, repetition, individual-word comprehension and semantic knowledge should be separated conceptually. The syndrome label follows the dominant deficit and its progression; disease pathology should not be inferred with more certainty than the evidence permits. Hearing, education, multilingual usage and premorbid vocabulary can distort a superficial language examination.

STEM PATTERN	BEST OPENING FORMULATION	DIAGNOSTIC PIVOT	CLASSIC WRONG TURN
Gradual memory-led loss with functional transfer	Progressive amnestic neurocognitive syndrome	Learning and retention pattern linked to decline	Naming a pathology from "forgetfulness" alone
Fluctuation with visual, attentional and motor features	Cortical synucleinopathy dementia phenotype	Convergent feature cluster and medication sensitivity	Using an isolated perceptual symptom as proof
Executive decline with vascular chronology	Vascular dementia syndrome	Temporal and anatomical coherence	Equating any vascular scan change with cause
Early social and behavioural disintegration	Frontotemporal neurodegenerative syndrome	Progressive change from personality baseline	Choosing psychiatric or neurological causation by stereotype
Progressive language-led impairment	Language-predominant degenerative syndrome	Nature of the language breakdown	Calling every word-finding complaint aphasia
Cognitive complaints embedded in depressive illness	State-related cognitive impairment or coexisting decline	Longitudinal consistency and function	Treating "pseudodementia" as a final diagnosis
Rapid, discordant or medically unstable decline	Urgent cognitive syndrome of uncertain cause	Tempo and evidence of active brain dysfunction	Assuming rapid decline is merely severe chronic dementia

28.5 Vignette: depression, effort and apparent reversibility

A patient with pervasive low mood, loss of interest, sleep disturbance and marked self-criticism reports catastrophic memory failure. Answers are frequently abandoned, yet performance improves with encouragement and structure. Family members describe reduced initiative but disagree about whether established skills have actually been lost. The referral asks whether this is "pseudodementia."

The term **depression-associated cognitive impairment** can describe the clinical problem, but it must not end the reasoning. Depression can impair attention, processing efficiency, retrieval and motivation; neurocognitive disease can present with depression; both can coexist. Distress about failure, variable engagement or ready use of "I do not know" is suggestive only in context. These are not stand-alone discriminators.

The answer should reconstruct premorbid function, establish chronology, seek objective and ecologically valid decline, examine the internal consistency of errors, treat the depressive state and follow cognition longitudinally. Improvement in mood with persistent cognitive or functional decline should reopen the etiological formulation. Improvement in both supports a substantial state-related contribution but does not erase vulnerability or eliminate the need for follow-up when concern remains.

PITFALL

Do not reassure the family that cognition is "only depression" after an initial interview. That phrase can delay recognition of coexisting neurocognitive disease. Record the uncertainty, treat what is treatable, define the functional observations to be followed and obtain collateral information again.

A related stem may describe gait deterioration, urinary symptoms, cognitive slowing and ventricular enlargement, or cognitive decline associated with systemic illness, medication burden, sensory loss or nutritional risk. These patterns should trigger a structured search for treatable contributors rather than a promise of complete reversibility. The endocrine, metabolic and nutritional diseases themselves belong in the Endocrine and metabolic psychiatry notes. The dementia answer owns the workup logic: identify a plausible mechanism, test it appropriately, treat it and reassess what remains.

28.6 Vignette: rapid decline and diagnostic urgency

A previously independent adult develops rapidly worsening cognition, behavioural change, gait instability and abnormal involuntary movements. The course is too fast for the common chronic degenerative pattern suggested in the referral. Mental state and neurological findings vary, and the patient is becoming medically difficult to manage at home. The examination asks for the approach rather than a rare diagnostic label.

The correct response is an **urgent rapidly progressive cognitive syndrome**. Tempo changes both differential weighting and locus of care. The candidate should confirm that the apparent rapidity is real by identifying the previous baseline, earliest unequivocal change and sequence of loss. A long unrecognised decline with recent caregiver collapse is different from genuinely rapid brain dysfunction. Delirium, seizures, toxic exposure, systemic disease, inflammatory or infectious processes and rapidly progressive degenerative disease must be considered according to the clinical context.

Assessment should become targeted and parallel rather than leisurely and serial. Stabilise immediate medical and behavioural risk, review substances and medications, obtain collateral chronology, perform full mental state and neurological examinations, and select laboratory, imaging, electrophysiological or fluid studies to answer specific hypotheses. Imaging physics belongs in the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes. The cerebrospinal-fluid examination procedure belongs in the Neurology foundations for psychiatrists notes. The autoimmune and limbic encephalitis workup belongs in the Neuropsychiatry of systemic and neurological disease notes.

The key is not to write an encyclopaedic list. Rank possibilities by tempo, systemic context, neurological pattern and treatability. State which diagnoses require immediate exclusion because delay changes outcome. Preserve infection-control, consent and capacity considerations where relevant, but do not invent statutory particulars. If a chronic diagnosis already appears in the notes, it must not be allowed to explain away a new, rapid or fluctuating deterioration.

28.7 Vignette: behaviour, prescribing and caregiver collapse

A person with established dementia begins pacing, shouting and striking out during evening care. The caregiver requests a sedating prescription and says home care may no longer be possible. The patient cannot give a coherent account, has reduced oral intake and appears uncomfortable during movement. The question asks for management.

LOCUS precedes treatment selection. Decide whether the person can be assessed safely in the present setting or needs urgent medical evaluation or a more supported environment. **FOCUS** is not "agitation" in the abstract. It is the immediate danger, distress and probable unmet need, together with caregiver exhaustion. **MODUS** begins with assessment for pain, acute illness, medication effects, sensory overload, communication failure, fear, hunger, toileting needs, sleep disruption and environmental triggers. Behavioural observation and caregiver reconstruction of antecedents and consequences are clinical data.

The **behavioural formulation** should describe what happens, when, with whom, during which task and with what apparent consequence. A label such as aggression compresses too much. If the behaviour is defensive during painful movement or intimate care, sedation may obscure the cause while increasing harm. Modify the task, environment and communication; treat identified contributors; support the caregiver; and define the target behaviour before considering disease-specific medication guidance.

- State the immediate risk to the patient, caregiver and others, and identify what would make the current setting unsafe.
- Look for a new medical, neurological, medication-related, environmental or interpersonal precipitant before calling the behaviour progression.
- Use simple communication, predictable care, sensory aids, meaningful activity and environmental modification matched to the trigger.
- If medication is considered, define the target, expected benefit, major harms, review plan and stopping logic using current dementia-specific guidance.
- Assess caregiver capacity and respite needs directly; "family support" is neither limitless nor clinically neutral.

IN INDIA

A workable plan must identify who actually provides supervision, who can accompany urgent assessment, who controls medicines and money, and what local services are reachable. Joint-family language can conceal that an exhausted relative carries nearly all care. Ask for the real roster and build the least restrictive safe plan around available people, transport, cost and follow-up.

Drug selection, dementia-specific dosing and adverse-event safeguards belong in the dedicated behavioural-management section. The candidate should never improvise a dose from general psychiatric practice. The answer earns marks by showing that medication is considered only after the target, cause, risk and non-drug response are defined, except where immediate danger requires simultaneous containment and assessment.

28.8 Write the final answer with VITALS

Under examination pressure, use a consistent retrieval spine. The aim is not to force different diseases into the same presentation. It is to ensure that the answer does not lose baseline, function, competing explanations or safety while pursuing a recognisable phenotype.

MNEMONIC

VITALS

Verify acquired decline from the person's baseline. **I**dentify the leading cognitive, behavioural and neurological pattern. **T**race tempo, sequence and functional consequence. **A**ssess alternatives, active contributors and acute threats. **L**ink collateral history, examination and targeted investigations into a ranked formulation. **S**tate the safest next step, uncertainty and follow-up plan.

For a diagnosis question, lead with the syndrome and its level of certainty, then justify the most likely cause using the pivot. For a differentiation question, compare the leading formulations on shared axes: onset, tempo, cognitive pattern, associated behaviour, neurological findings, functional trajectory and response to context. For an investigation question, state the hypothesis each test addresses. For a management question, use LOCUS, FOCUS and MODUS, then connect every intervention to a target and review plan.

The **discriminating negative** is an absent expected feature that genuinely weakens a competitor. Use it carefully. Absence is meaningful only when the feature was sought reliably and should be present at the relevant stage. Do not turn silence in a short stem into proof that a symptom is absent. Likewise, do not invent normal findings. State what information is missing and how it would change the ranking.

When investigations disagree with the phenotype, name the discordance. A structural abnormality may not explain the clinical distribution; a disease-associated biomarker may not explain function; a reassuring brief screen may not outweigh credible longitudinal decline. Recheck sampling, context, language, education, sensory function, comorbidity and mixed pathology. General test administration remains in the Psychometry, Assessment and Descriptive Psychopathology notes, while dementia-specific interpretation belongs in this chapter.

End with action. A complete formulation that does not change care is unfinished. Clarify immediate safety, the setting of care, reversible contributors, support for the caregiver, treatment targets, monitoring and what finding should prompt reassessment. This is also the correct response to uncertainty: not false precision, but a safer plan for acquiring better evidence.

In every dementia vignette, verify decline, identify the pattern and tempo, establish functional consequence, rank cause against its strongest competitor, and state the action that cannot safely wait.

29 · Consolidation and quick review

The marks lie in integration, not in reciting isolated lists. A strong answer to a cognitive-disorder question moves in a disciplined direction: establish that there is a cognitive syndrome, define its pattern and tempo, determine whether everyday function has changed, search for reversible or superimposed contributors, formulate the most defensible cause, and convert that formulation into a safe plan. The same sequence protects clinical practice. It prevents a familiar complaint such as forgetfulness from being mistaken for a diagnosis and prevents a familiar scan abnormality from replacing bedside reasoning.

Consolidation means learning a reusable architecture rather than another catalogue. Whether the doorway is a memory complaint, behavioural change, executive failure, an imaging finding or a caregiver's safety question, return to the same questions. What syndrome is present? How secure is the evidence? What explains it? What contributes? What matters now?

29.1 Build the syndrome before naming the disease

The first task is to distinguish a cognitive complaint from a demonstrable **cognitive syndrome**. Subjective concern deserves assessment, but it does not by itself establish objective impairment. Conversely, a patient who minimises difficulty may still have substantial change recognised by family or colleagues. The clinician integrates self-report, collateral account, observed performance and change from the person's previous level. Premorbid ability, education, language, occupation, sensory impairment and opportunity to perform a task all shape what a test result means.

Next determine whether cognition has altered everyday independence. **Functional decline** is not simply inability to perform a task during an interview. It is a change in real-world execution attributable to cognitive difficulty rather than physical illness, poverty, unfamiliarity, lack of opportunity or another psychiatric state. Complex activities often reveal change before basic self-care, but the important issue is the causal link between cognition and function. Ask what the person previously did, what now goes wrong, who compensates, and what would happen without that compensation.

Syndrome	Function	Collateral	Safety
BEFORE SUBTYPE	BEFORE LABEL	BEFORE CERTAINTY	BEFORE ELEGANCE

The syndrome statement should be compact but rich: the dominant cognitive pattern, associated behavioural or neurological change, tempo, functional consequence and any active confounder. Etiology comes after this statement. A disease label given too early creates confirmation bias; subsequent history is then forced to fit it. A syndrome-first formulation leaves room for mixed pathology, treatable contributors and diagnostic revision.

RULE **Name the level and pattern of impairment before naming its cause.** A defensible etiological diagnosis rests on a demonstrated syndrome, meaningful functional evidence and a trajectory that coheres with the proposed explanation.

29.2 Make history longitudinal and collateral

A cross-sectional interview can show how the patient performs today; dementia assessment asks how the person has changed over time. The core evidence is therefore longitudinal. Establish the earliest credible change, the sequence in which difficulties emerged, whether progression has been smooth or stepwise, and whether fluctuations represent genuine variation, environmental demand, fatigue, sleep disruption, medication effects or an acute confusional state. Tempo narrows the differential far more efficiently than an undirected list of causes.

Collateral history is not an optional confirmation of the patient's account. It is an independent source of evidence about decline, compensation and risk. Speak to someone who knows the patient's previous abilities and current routine. Resolve disagreements rather than automatically privileging either narrator. A patient may under-report because of reduced awareness, shame or fear of lost autonomy. A relative may over-report because of anxiety, conflict, burden or unrealistic expectations. The clinician's task is triangulation.

- Map change across memory, language, attention, visuospatial ability, executive control, social cognition and behaviour without assuming that memory must dominate.
- Translate each reported deficit into a concrete event such as missed payments, repeated journeys, unsafe cooking, lost work sequences, altered social judgement or failure to use familiar objects; ask who has silently taken over tasks, because family compensation can conceal clinically important loss of independence.
- Review medical illness, prescribed and non-prescribed substances, sleep, pain, mood, psychosis, sensory loss and environmental disruption as contributors; record driving, medication management, finances, falls, wandering, self-neglect, vulnerability to exploitation and access to weapons or toxic substances as safety domains, not as moral judgements.

The history should also identify the patient's values and the family's explanatory model. Some families present behavioural change as stubbornness, laziness or normal ageing; others interpret ordinary lapses as inevitable dementia. Correcting either error requires examples and functional reasoning, not reassurance alone. Language concordance and culturally familiar tasks improve the validity of the encounter, but cultural sensitivity must not become an excuse for avoiding a clear assessment of impairment.

29.3 Use examination and investigations as converging evidence

The mental state examination, bedside cognition and neurological examination answer different questions. The interview reveals engagement, affect, thought content, perception, insight and behaviour. Cognitive examination samples domains and exposes error patterns. Neurological examination looks for focal, extrapyramidal, cerebellar, gait, motor, sensory or release phenomena that may redirect the formulation. No single component should be treated as a universal arbiter.

The **cognitive phenotype** is the pattern formed by relative strengths, weaknesses and characteristic errors. A total screening score compresses that pattern and can be useful for communication or follow-up, but it cannot explain why the score is low. Performance may be distorted by education, language, poor hearing or vision, anxiety, depression, pain, fatigue and

reduced effort. General administration and scoring of standardised screens belong in the Psychometry, Assessment and Descriptive Psychopathology notes. In a dementia assessment, the relevant question is how screening findings fit the history, function and observed behaviour.

EVIDENCE LAYER	QUESTION IT ANSWERS	COMMON MISUSE	CORRECT INTEGRATION
Patient account	What is experienced and feared?	Treating confidence as accuracy	Compare with previous ability and other evidence
Collateral account	What has changed in daily life?	Accepting family interpretation as fact	Seek examples, sequence, compensation and context
Mental state	What psychiatric state affects cognition?	Calling all poor performance dementia	Separate state effects from enduring decline
Cognitive examination	Which domains and error types are affected?	Using a total score as an etiology	Interpret the profile with language and education
Neurological examination	What accompanying brain-system signs are present?	Ignoring subtle signs after a psychiatric presentation	Use signs to refine localization and differential diagnosis
Laboratory assessment	Is a contributor or mimic present?	Ordering without a clinical question	Link each test to a plausible, actionable hypothesis
Imaging	Is there structural or functional support?	Equating an abnormality with causation	Ask whether distribution and severity fit the syndrome
Fluid biomarker	Does biology support a proposed pathology?	Allowing biology to erase the phenotype	Update diagnostic probability and management eligibility

The physics and general interpretation of imaging modalities belong in the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes. The cerebrospinal-fluid examination procedure belongs in the Neurology foundations for psychiatrists notes. Here, their dementia-specific purpose is conceptual: investigations test a hypothesis generated at the bedside, exclude important alternatives, and sometimes increase confidence in an underlying pathology. They support a clinical formulation; they do not manufacture one.

29.4 Protect the differential from premature closure

The differential diagnosis is most efficient when organised by tempo, cognitive pattern, associated features and reversibility rather than by a memorised undifferentiated list. Begin by asking whether the presentation is chronic and progressive, acute or fluctuating, rapidly evolving, stepwise, static, or apparently variable with mood and effort. Then ask whether the dominant deficit is amnestic, executive, language-led, visuospatial, behavioural, subcortical or mixed. Finally, identify systemic, neurological, psychiatric, medication-related and environmental explanations that could mimic or worsen the syndrome.

Reversible contributors should be framed carefully. Finding a potentially treatable abnormality does not prove that it explains the whole cognitive syndrome, and treating it does not guarantee restoration of previous cognition. The clinically useful formulation often contains layers: an underlying degenerative or vascular process, a superimposed medical or psychiatric burden, and modifiable factors that amplify disability. This layered model avoids the false choice between "reversible" and "irreversible." The endocrine, metabolic and nutritional diseases themselves are covered in the Endocrine and metabolic psychiatry notes.

Rapidly progressive decline requires urgency because the diagnostic balance shifts toward active neurological, inflammatory, infectious, toxic, metabolic and other potentially actionable processes. The detailed workup of autoimmune and limbic encephalitis belongs in the Neuropsychiatry of systemic and neurological disease notes. Delirium is taught in the Stroke, TBI, delirium and focal brain syndromes notes. The dementia clinician must nevertheless recognise when either possibility changes the locus and urgency of assessment.

Depression can produce prominent cognitive symptoms, and a neurocognitive disorder can coexist with depression. Variable engagement, subjective distress or apparent effort do not settle the question. The distinction depends on longitudinal evidence, internal consistency, functional examples, mental state and response over time. Similarly, vascular brain injury may contribute to a dementia syndrome, while the broader pre-dementia and post-stroke cognitive spectrum belongs in the Stroke, TBI, delirium and focal brain syndromes notes.

TRAP **Do not answer "dementia" when the vignette asks for a cause, and do not answer with a cause before establishing dementia.** Examiners reward the chain from syndrome to phenotype to etiology, with competing explanations explicitly addressed.

29.5 Formulate certainty, mixtures and discordance

A useful diagnostic formulation states not only what is most likely, but why, how certain the clinician is, and what evidence does not fit. The level of cognitive disorder, dominant phenotype, probable cause, contributing conditions, behavioural complications and functional risks should be distinguishable. This structure allows the diagnosis to evolve without making earlier work appear careless. Neurocognitive disorders are often probabilistic clinical constructs, especially when several pathologies or systemic burdens plausibly coexist.

Mixed etiology should not become a refuge for uncertainty. It is appropriate when independent lines of evidence support more than one meaningful contributor. It is weak when used merely because the history is incomplete or the scan contains common abnormalities. State which contributor explains which aspect of the presentation, which evidence is shared, and what future observation could shift the balance.

A **biomarker** changes the probability of a biological process; it does not independently establish the patient's functional state or explain every symptom. Concordance among phenotype, trajectory, examination and biomarker strengthens a formulation. Discordance should trigger a deliberate review of sampling, test limitations, comorbidity, disease stage and the possibility that

more than one process is present. It should not trigger selective forgetting of inconvenient clinical evidence.

GOOD TO REMEMBER

Write the formulation as a hierarchy: syndrome and severity of functional effect; dominant cognitive and behavioural pattern; most likely etiology; contributors and mimics; current risks; and the evidence still needed. This is clinically useful and converts directly into an examination answer.

Diagnostic language should preserve uncertainty without becoming vague. Terms expressing probability are useful only when the evidence behind them is stated. Naming discordant data is mature reasoning. The goal is not maximal certainty; it is maximal defensibility.

29.6 Translate formulation into a whole-care plan

Management begins with **LOCUS**: decide whether assessment and care can proceed in the outpatient setting or whether acute risk, rapid deterioration, severe behavioural disturbance, inability to maintain basic needs, caregiver collapse or diagnostic urgency requires a more supported setting. Locus is dynamic. It changes when risk, supervision, medical stability or available support changes.

FOCUS identifies the targets of care across phases. Immediate priorities are safety, distress, acute medical or psychiatric contributors and communication of the working diagnosis. Subsequent priorities include preserving function, reducing avoidable disability, managing cognitive and behavioural symptoms, supporting the caregiver, planning supervision, and revisiting goals as the illness changes. The target must be explicit. “Treat dementia” is not a treatment target; unsafe wandering, distressing misidentification, medication errors, insomnia, caregiver exhaustion and loss of meaningful activity are.

MODUS ensures that the plan is not reduced to a prescription. Depending on the formulation, modes include disease-specific medication, treatment of contributors, environmental adaptation, psychological strategies, occupational support, rehabilitation, caregiver education, social and legal planning, and palliative approaches. Cognitive enhancers, behavioural and psychological symptoms, antipsychotic risk and disease-modifying anti-amyloid therapy are taught in their dedicated sections of this chapter. General antipsychotic pharmacology belongs in the Schizophrenia: management, antipsychotics and adverse effects notes.

Every intervention needs a reason, a realistic outcome, a monitoring plan and a stopping rule. Benefits should be described in functional or behavioural terms that matter to the patient and family. Risk discussion should include uncertainty and burden, especially when treatment requires repeated monitoring or when the person cannot provide fully informed consent. Capacity is decision-specific and should be assessed in relation to the actual choice rather than inferred globally from the diagnosis.

IN INDIA

Plans must fit the family's actual geography, finances, language, caregiving structure and access to investigations, rehabilitation and follow-up. A theoretically complete plan that cannot be implemented is not superior to a staged plan that addresses danger first, uses available supports and preserves continuity. Ask who will supervise medication, accompany visits, manage money and provide respite; do not write “family support” as if it were an unlimited resource.

29.7 Follow the person, not merely the score

Dementia care is longitudinal. Follow-up should revisit cognition, function, behaviour, neurological change, medical contributors, medication benefit and harm, nutrition, mobility, sleep, continence, pain, sensory function and safety. The relative importance of these domains changes over time. A stable screening score can coexist with meaningful functional deterioration, while a lower score may reflect language, illness, fatigue or testing conditions rather than a true change in trajectory.

The **care dyad** is the patient and the person providing sustained support. Their needs overlap but are not identical. The patient's autonomy, comfort, identity and preferences remain central; the caregiver's health, sleep, burden, knowledge, grief and ability to continue must be assessed directly. A plan that protects the patient by exhausting the caregiver is unstable. Equally, caregiver convenience cannot automatically override the patient's rights and preferences.

- Define what change would count as benefit, deterioration or unacceptable harm before altering treatment, and review adherence and administration rather than assuming that prescribed treatment is received correctly.
- Reassess environmental triggers and unmet needs before attributing every new behaviour to disease progression; update supervision, financial safeguards, mobility support and crisis plans as function changes.
- Revisit communication preferences, future care goals and the caregiver's ability to sustain the plan.

PITFALL

A new behavioural or cognitive deterioration should not be casually labelled “progression.” First look for pain, infection, constipation, medication effects, sleep disruption, sensory loss, environmental change, depression and delirium. The label of dementia does not protect a patient from developing another illness.

End-of-life care is not an abrupt final chapter of management. It grows from earlier conversations about values, burdens, comfort, feeding, hospital transfer and acceptable intervention. These discussions require tact, repetition and attention to capacity and family conflict. The therapeutic task is to preserve personhood while making progressively realistic decisions.

29.8 The master retrieval frame

When the vignette feels crowded, return to a fixed retrieval frame. The aim is not to force every case into a slogan. It is to prevent omission under pressure. The frame below follows the natural order of assessment and care, while leaving disease-specific criteria, investigations and treatments to be inserted only when the case supports them.

MNEMONIC

MIND CARE

Map the complaint, baseline and tempo. **I**dentify the cognitive and behavioural pattern. **N**ote functional change, neurological signs and safety. **D**etect delirium, depression, drugs and reversible contributors. **C**orroborate with collateral history, examination and targeted investigations. **A**ssign the syndrome, likely etiology and degree of certainty. **R**espond through LOCUS, FOCUS and MODUS for the patient and caregiver. **E**valuate longitudinally and revise the formulation when evidence changes.

In a long answer, begin with a brief definition of the syndrome in clinical terms, then classify the presentation by level, phenotype, tempo and likely cause. Follow with history, collateral information, mental state, cognition, neurological examination and targeted investigations. Present the differential as a reasoned hierarchy. End with a whole-care plan and follow-up parameters. In a short answer or viva, compress the same sequence rather than substituting an unrelated list.

The boundaries around neighbouring topics should remain visible. General cognitive-screen administration belongs in the Psychometry, Assessment and Descriptive Psychopathology notes. Imaging physics belongs in the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes. The cerebrospinal-fluid procedure belongs in the Neurology foundations for psychiatrists notes. Delirium belongs in the Stroke, TBI, delirium and focal brain syndromes notes. Movement-disorder neuropsychiatry, inherited movement disorders and autoimmune encephalitis belong in the Neuropsychiatry of systemic and neurological disease notes. Intellectual disability as a developmental condition belongs in the Intellectual disability and autism spectrum disorder notes. Cross-reference these areas without importing their criteria into a dementia answer.

A final answer should sound like clinical reasoning: evidence first, inference second, action third. State what establishes the syndrome, supports the etiology, argues against alternatives, remains modifiable and matters now. If evidence is incomplete, say what remains uncertain and how it would be resolved.

The must-memorise sequence is: establish syndrome, map phenotype and tempo, verify function with collateral evidence, exclude superimposed and treatable contributors, formulate cause and certainty, then plan care for both patient and caregiver.

References

The canonical sources this chapter is grounded in. Every fact-bearing specific in the notes traces to one of these; where a primary paper or a trial is named in the text, it is cited inline. The list is ordered by how much of the chapter rests on each.

1. Boland R, Verduin ML, eds. Kaplan and Sadock's Comprehensive Textbook of Psychiatry. 11th ed. Wolters Kluwer; 2024.
2. Geddes JR, Andreasen NC, Goodwin GM, eds. New Oxford Textbook of Psychiatry. 3rd ed. Oxford University Press.
3. Taylor DM, Barnes TRE, Young AH. The Maudsley Prescribing Guidelines in Psychiatry. 15th ed. Wiley-Blackwell.
4. American Psychiatric Association. Diagnostic and Statistical Manual of Mental Disorders, 5th ed, Text Revision (DSM-5-TR). APA; 2022.
5. Kaufman DM, Geyer HL, Milstein MJ. Kaufman's Clinical Neurology for Psychiatrists. 9th ed. Elsevier.
6. Johnstone EC, Owens DGC, Lawrie SM, et al, eds. Companion to Psychiatric Studies. 8th ed. Churchill Livingstone/Elsevier.
7. Hodges JR. Cognitive Assessment for Clinicians. 3rd ed. Oxford University Press.
8. World Health Organization. ICD-11 Clinical Descriptions and Diagnostic Requirements for Mental, Behavioural and Neurodevelopmental Disorders (CDDR). WHO; 2024.
9. Stahl SM. Stahl's Essential Psychopharmacology: Neuroscientific Basis and Practical Applications. 5th ed. Cambridge University Press.
10. Tasman A, Riba MB, Alarcon RD, et al, eds. Tasman's Psychiatry. 5th ed. Springer.
11. Brunton LL, Knollmann BC, eds. Goodman and Gilman's The Pharmacological Basis of Therapeutics. McGraw-Hill.
12. McKeith IG, Boeve BF, Dickson DW, et al. Diagnosis and management of dementia with Lewy bodies: fourth consensus report of the DLB Consortium. Neurology 2017;89(1):88-100. Reproduced as a diagnostic-criteria standard with acknowledgement.
13. Hachinski VC, Iliff LD, Zilhka E, et al. Cerebral blood flow in dementia. Arch Neurol 1975;32(9):632-637. The ischaemic score is described, not reproduced as a scored table.
14. van Dyck CH, Swanson CJ, Aisen P, et al. Lecanemab in early Alzheimer's disease (Clarity-AD). N Engl J Med 2023;388(1):9-21.
15. Sims JR, Zimmer JA, Evans CD, et al. Donanemab in early symptomatic Alzheimer disease (TRAILBLAZER-ALZ 2). JAMA 2023;330(6):512-527.

Additional sources cited inline where a specific rests on them: the Prescriber's Guide (Stahl), the Geriatric Psychiatry study guide of the American Psychiatric Association, and 50 Studies Every Psychiatrist Should Know. The MMSE and MoCA are named as instruments and described in our own words; no item content is reproduced (both are copyright-restricted). Guideline positions are attributed in the text to the body that issues them, and where a guideline position and Indian practice diverge, both are named and neither is presented as the other.

PART IV

IV

Endocrine and Metabolic Psychiatry

When the gland is driving the mind: thyroid, adrenal, glucose and the metabolic cost of treatment, read in both directions.

01 The endocrine axes: adrenal, thyroid and parathyroid

02 Metabolic and nutritional disease

03 The metabolic cost of treatment, diabetes and electrolytes

04 Encephalopathies, reproductive endocrinology and toxins

05 Consolidate and apply

VISUAL 4.1



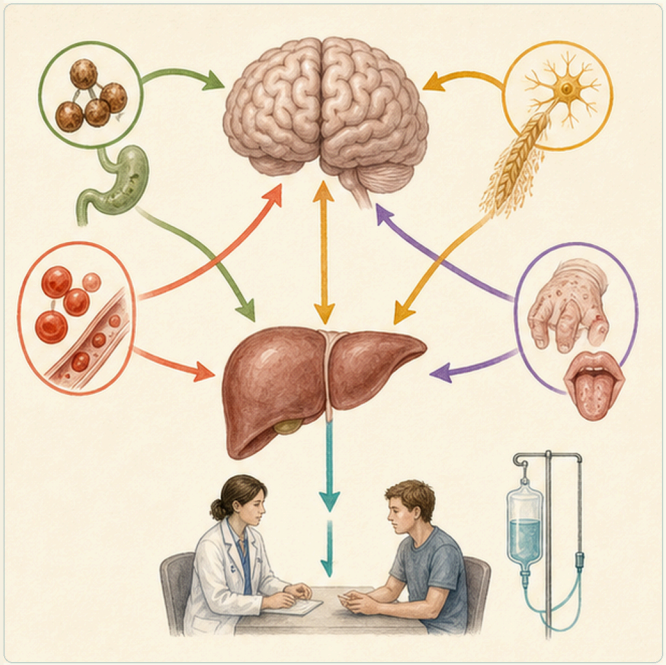
Endocrine axes produce psychiatric phenotypes

- HPA and adrenal
- Thyroid
- Parathyroid and calcium
- Psychiatric phenotype
- Laboratory context

Endocrine disease can alter mood, cognition, arousal and behaviour; diagnosis requires the medical pattern and laboratory context.

Source lock: day-33.final.html#sec-t01

VISUAL 4.2



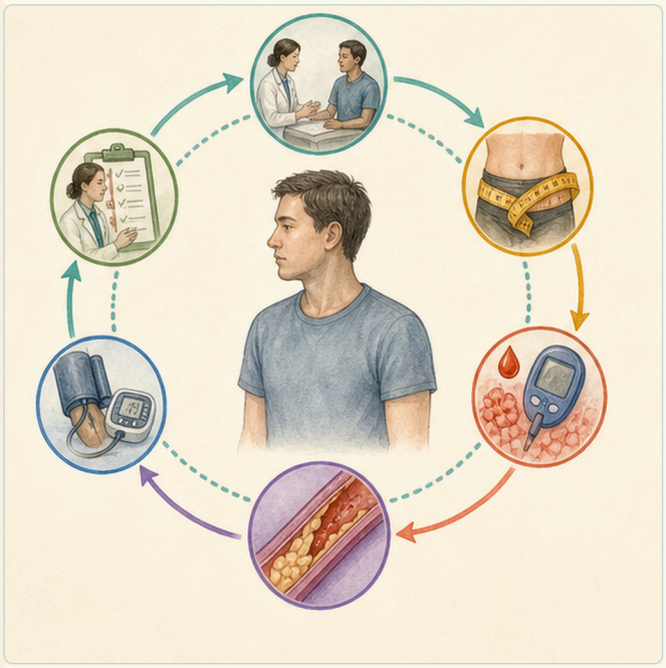
Metabolic and nutritional syndromes leave signatures

- Copper
- Thiamine
- B12 and folate
- Niacin
- Treat the cause

Metabolic and nutritional syndromes are recognised through their systemic and neurological signatures and treated at the cause.

Source lock: day-33.final.html#sec-t14

VISUAL 4.3



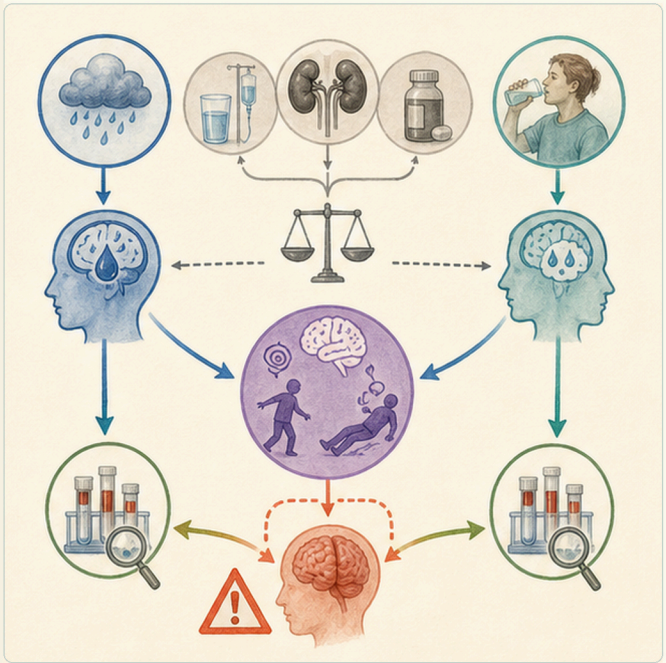
Metabolic monitoring is a repeated clinical loop

- Baseline and review
- Weight and waist
- Glucose
- Lipids
- Blood pressure
- Intervene

Metabolic monitoring is useful only when repeated results lead to lifestyle, medical or prescribing action.

Source lock: day-33.final.html#sec-t16

VISUAL 4.4



Sodium disorders require cause and correction control

- SIADH
- Polydipsia
- Volume and medication context
- Neurological symptoms
- Correction hazard
- Serial monitoring

Hyponatraemia management depends on cause, chronicity, symptoms and controlled correction with serial monitoring.

Source lock: day-33.final.html#sec-t21

PAPER IV

Endocrine and metabolic psychiatry

Four bands . one complete notes document

The endocrine axes: adrenal, thyroid and parathyroid

1 · Endocrine and metabolic disease for the psychiatrist: what this chapter owns, and what it points to

Every psychiatric syndrome in this chapter can be produced by a gland, a nutrient, a solute, a metal or a drug that a psychiatrist prescribed. The examinable point is not the soft one that hormones influence mood. It is the hard one: a real proportion of psychiatric presentations are generated, sustained or amplified by a disturbance of an endocrine axis, a metabolic pathway, a nutrient store, an electrolyte, an organ of clearance or an environmental toxin, and a large share of those disturbances are correctable. Depression, anxiety, psychosis, catatonia, apathy, cognitive decline and personality change all recur across this chapter, arriving each time from a different organ. The mental state examination alone will not say which organ, and that single limitation is the problem the chapter exists to solve.

The work is therefore not memorising a catalogue of glands and their psychiatric associations. It is learning to hold an ordinary psychiatric presentation and ask a disciplined sequence of questions about the body that produced it: which system could generate this picture, which way causation runs, what in the history or examination points to an organ rather than a mood disorder, which test answers that question, and what changes if it comes back abnormal. The disease sections supply the criteria, investigations, thresholds and treatments. This orientation supplies the **clinical architecture** they sit inside, and the boundary map of what belongs here and what belongs to a neighbouring chapter.

1.1 Traffic runs both ways between body and mind

Most candidates arrive with one model: systemic disease produces psychiatric symptoms. That covers part of the chapter and misreads the rest. The organising discipline is to establish the **direction of causation** first, because it determines who investigates, who treats, and what a good answer looks like.

Systemic disease may present to psychiatry wearing a psychiatric face, referred as depression or psychosis while the disease stays silent, subtle or attributed to the mental illness. Causation also runs the other way, where psychiatric illness and its treatment generate endocrine and metabolic disease: the exposure is a prescription and the outcome is measured in weight, lipids, glucose, prolactin and sodium. The two may also coexist on shared ground, because severe mental illness, poverty, inactivity, diet, tobacco and reduced access to physical healthcare all move metabolic risk the same way, so association establishes nothing about cause. Finally the treatments collide: a psychotropic unsafe in organ failure, a hormone replacement that unmasks an affective episode, an endocrine therapy that alters psychotropic handling.

DIRECTION OF CAUSATION	HOW THE CASE USUALLY ENTERS	THE REASONING IT DEMANDS	WHERE DEPTH FOLLOWS
Systemic disease presenting psychiatrically	Referral for a mood, psychotic or cognitive syndrome	Find the somatic signature; treat the disease, then reassess the mental state	The endocrine axis, nutritional and encephalopathy sections
Psychiatric treatment producing systemic disease	An abnormal weight, waist, lipid, glucose, prolactin or sodium result	Attribute to drug, illness or baseline; switch, add or accept	The treatment-burden and monitoring sections
Shared ground and co-occurrence	A patient with both, and no clear sequence	Resist causal language; treat both, explain neither by the other	The diabetes, obesity and mortality sections
Treatment collision	A prescribing question in organ impairment or hormone therapy	Predict the interaction; choose the agent the failing organ tolerates	The prescribing-in-impairment and reproductive sections

RULE **Name the direction of causation before you name the disease.** The same laboratory abnormality means something different depending on whether the body generated the psychiatric state, the treatment generated the abnormality, or both merely travelled together. Answers that leave the direction unstated read as association-hunting and collapse under one follow-up question in the viva.

1.2 Reversibility is the currency of this chapter

These diseases earn their disproportionate examination weight because of what happens when they are found. Correct the deficiency, replace the hormone, remove the toxin, chelate the metal, restore the solute, and a psychiatric syndrome that looked autonomous may remit, sometimes completely. That converts a screening decision into a clinical duty, and it is why the chapter reads as a set of things not to miss rather than a survey of endocrinology.

Reversibility is not unconditional, and the qualification is where the marks sit. Several diseases here have a **treatment window**: an interval during which correction restores function, beyond which a residual deficit becomes fixed. Others reverse only partially, leaving a remnant to be recognised as sequel rather than pursued as fresh pathology. Others cause harm through the correction itself, so the rate of treatment matters as much as the fact of it. Each belongs to the section that owns it; the rule here is that reversibility is a claim with conditions attached, never a promise.

Axis	Direction	Window	Boundary
THE SYSTEM	THE CAUSAL RUN	THE URGENCY	THE HANDOVER

This governs how cognitive presentations are handled. When a metabolic, endocrine or nutritional disease produces cognitive impairment, this chapter teaches the disease, its cognitive signature and the extent to which it reverses; the dementia syndrome and its workup belong to the Dementias and neurocognitive disorders notes. When systemic disturbance produces an acute confusional state, this chapter teaches the disease and the derangement that drove it, while

delirium as a syndrome, with its criteria, assessment and prevention, belongs to the Stroke, TBI, delirium and focal brain syndromes notes. Keeping that division clean prevents the commonest structural fault here, which is spending half the page rewriting a workup the question did not ask for.

1.3 The mental state cannot name the organ; the body can

Because the psychiatric phenotypes overlap so heavily across the chapter, the discriminating evidence is almost never in the mental state. It is in the **somatic signature**: the physical symptoms, examination findings, temporal features and exposures that point at one organ rather than another. Depression with cold intolerance, hair and skin change and slowed relaxation of reflexes is a different proposition from depression with weight loss, tremor, heat intolerance and tachycardia, and both differ from depression with pigmentation, postural dizziness and salt craving. The affective content may be indistinguishable in all three.

This puts an old-fashioned demand on the psychiatric history and examination. A systematic review of systems is not padding; it is the discriminating instrument. So is a drug history covering prescribed steroids, over-the-counter preparations, traditional and complementary remedies, supplements and anything the patient does not consider a medicine, and an exposure history covering occupation, water source, hobby chemicals, home remedies and cosmetics, alongside nutritional, alcohol, menstrual and reproductive histories and a family history that may carry an inherited metabolic disorder.

Certain features should raise suspicion sharply and are worth carrying as a standing list, because they are the practical trigger for looking beyond the psychiatric diagnosis.

- Onset atypical for the psychiatric diagnosis by age, tempo or absence of a plausible psychosocial context.
- A mixed or shifting picture crossing diagnostic boundaries, especially affective symptoms with psychosis, cognitive change or fluctuating conscious level.
- Prominent physical symptoms the patient has not connected to the psychiatric presentation.
- Failure to respond to adequate psychiatric treatment, or unexpected deterioration during it.
- Neurological signs, movement abnormality, seizures or autonomic instability alongside the mental state.
- A drug, occupational or nutritional exposure with a plausible temporal relationship to onset.

PITFALL

The physical examination that was never done because the patient came through a psychiatric door. Pigmentation, a goitre, a tremor, an abnormal eye sign, a postural drop, a neuropathy or an enlarged liver will not appear in the notes if nobody looked, and the formulation then hardens around an absence of evidence the examination itself manufactured. Record what you found, not what you asked.

1.4 Investigate by question, not by panel

The reflex is to order a broad screen, and it is only half right. Investigation here earns marks when each test answers a question the clinician has already formed. Two ideas are routinely merged: **screening** applies a test to an unselected group because the disease is common enough and missing it serious enough, while **case-finding** applies it to this patient because something in the history or examination raised its probability. The threshold to test is far lower in the second, and a borderline result means something different.

That interpretive difference is the chapter's most reliable trap. A mildly abnormal result in an unselected patient carries a far weaker claim on the diagnosis than the same value in a patient whose history already pointed at that organ. The direction of the abnormality, the pattern across related analytes, the clinical state at sampling, concurrent drugs and the acute illness itself all shift what a number means. The specific interpretations belong to the sections that own them: thyroid function tests to the hypothyroidism and hyperthyroidism section, biochemical confirmation of copper disorder to the Wilson's disease investigation section, the biomarkers probing the stress axis to the Cushing and steroid-psychiatry section.

TRAP **Declaring a borderline abnormality to be the cause of the psychiatric presentation.** Candidates find one deranged value, name it as the explanation and stop. A result at the edge of a reference interval in a patient with no supporting somatic signature is weak evidence, and acute illness of almost any kind perturbs several of these systems transiently. State instead what the result raises, what would confirm or refute it, and what the psychiatric management is meanwhile.

IN INDIA

Investigation strategy here is shaped by who pays and how far the patient travelled. Much of this testing is out of pocket, and a patient who has lost a day of wages to attend is unlikely to return to complete a staged workup. Settle the full question before the patient leaves the room, and be honest about which tests will change management rather than ordering a panel that reassures the clinician. Iodine status, endemic goitre and nutritional deficiency vary sharply by region and diet, so local prevalence should inform the pre-test probability rather than a figure recalled from an imported textbook. Steroids, prescribed and bought without prescription alike, are widely available, which makes the drug history a diagnostic instrument rather than a formality. Where specialist assays or an endocrinology opinion are distant, the psychiatric plan must be safe in the interval, not suspended awaiting a result.

1.5 The iatrogenic axis: what psychiatric treatment does to metabolism

Causation running from treatment to body inverts the clinical posture: here the psychiatrist is not the detective but the exposure. Antipsychotic treatment carries weight, lipid, glucose and prolactin consequences, mood stabilisers and antidepressants carry their own endocrine and electrolyte liabilities, and the resulting **iatrogenic burden** contributes materially to the excess physical morbidity and shortened life expectancy of people with severe mental illness. This is not incidental to psychiatric practice. It is one of its principal harms, and it is prescribed.

The chapter divides the territory carefully, and so should an answer. The receptor-mediated mechanism by which antipsychotic treatment produces metabolic effect is taught in the treatment-burden section, with the comparative liability of individual agents. The monitoring schedule and its cadence are stated once, in the monitoring section; the criteria and thresholds defining metabolic syndrome once, in the metabolic syndrome and mortality section; and management of established metabolic disturbance and raised prolactin in its own section. Drug selection and the receptor pharmacology underlying it belong to the Schizophrenia: management, antipsychotics and adverse effects notes; lithium as a drug, with its pharmacokinetics, therapeutic monitoring and non-thyroid organ effects, belongs to the Mood disorders: pharmacotherapy notes. This chapter takes only the endocrine and metabolic lens on each.

GOOD TO REMEMBER

Metabolic harm from psychotropic treatment is largely front-loaded, so fix the baseline before the first dose rather than discovering a change later without a comparator. Record weight and waist, ask about family history of diabetes and cardiovascular disease, and put the physical starting point in the case record. A patient whose baseline was never documented cannot be shown to have deteriorated, which disadvantages the patient far more than the audit.

1.6 Know what this chapter hands over

A chapter this broad touches many constructs it does not own, and an answer that re-teaches them is both long and wrong for the question. The **boundary map** is worth learning as content in its own right. Beyond the handovers already named, three more matter. The persistent amnestic syndrome that may follow acute thiamine deficiency is taught in the Stroke, TBI, delirium and focal brain syndromes notes; this chapter teaches the acute disease and its treatment window, then hands over the sequel. The basic biology of hormone synthesis, receptor signalling and neuroendocrine measurement is taught in the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes. The administration and scoring of cognitive and psychopathological instruments are taught in the Psychotherapies, clinical assessment, MSE and nosology notes, so cognitive change is described here in clinical language rather than through reproduced test items.

The habit that keeps the handover clean is to name the neighbouring construct in a clause and move on rather than pausing to define it, while saying explicitly where that chapter would change management. The examiner is testing whether you know the interface exists.

1.7 Managing the patient who has both

Management here is almost always joint, and it becomes coherent when built on the standard spine. **LOCUS** asks where care should happen. Most of this work is outpatient and shared with a physician or endocrinologist, but several presentations are genuine emergencies in which the derangement, not the psychiatric syndrome, dictates management on a medical ward or in critical care with psychiatric input rather than the reverse. Deciding LOCUS early prevents a common

failure: admitting an acutely deranged patient to a psychiatric bed that cannot deliver the monitoring the body needs.

FOCUS asks what is being treated and in what order. Acutely, the target is the derangement that threatens the patient, with enough containment of behaviour and risk to make treatment possible. In the short term it is the underlying disease and the syndrome it drove, on the expectation that the second may improve as the first is corrected. In the medium term it is the residual psychiatric picture persisting after correction, where independent psychiatric treatment becomes justified rather than assumed. In the long term it is relapse prevention, surveillance for recurrence of the systemic disease, and the cumulative physical cost of psychiatric treatment itself.

MODUS asks which modes are available and which are in use. The pharmacological mode covers both the treatment of the systemic disease and the psychotropic prescription, chosen for the organ that must clear it. The psychological and behavioural mode covers adherence support, dietary and activity intervention, tobacco and alcohol work, and psychoeducation the patient and family can act on. The procedural mode covers the surgical, endoscopic and replacement interventions the disease sections describe. The social mode covers the cost of long-term physical treatment, the accessibility of monitoring, family supervision of a complex regimen, and the exposure that may need to change before the illness will. Dosing, eligibility and monitoring intervals belong to the management sections and to current guidelines, not to memory.

1.8 Using the map under examination pressure

When the stem is crowded with physical detail, frame the case before you try to solve it. The scaffold below is a retrieval aid rather than a diagnostic criterion, and it asserts nothing this chapter does not teach. It holds the order of reasoning when the temptation is to seize the first recognisable physical sign and write everything known about the disease it suggests.

MNEMONIC

TRACE

Tempo and trigger: when the mental state changed, and what changed around it, including every drug started or stopped. **R**eversibility: is a treatment window open, and does it make investigation urgent today. **A**xis or organ: which endocrine axis, metabolic pathway, nutrient, solute or organ of clearance could generate this picture. **C**ausal direction: body driving mind, treatment driving body, the two merely travelling together, or treatments colliding. **E**vidence and escalation: which test answers the question actually asked, what would confirm or refute it, and when this patient needs a physician rather than another psychiatric review.

In a long answer, state the construct plainly, then move through the direction of causation, the somatic signature that identified the system, the targeted investigation and its interpretation, the differential within that system, and management by LOCUS, FOCUS and MODUS. In a short note the marks sit in the discriminating features and the test that separates the candidates, so lead with those rather than epidemiology. In a vignette the physical detail is placed deliberately;

read each item as evidence for or against an organ. In a viva, make the reasoning audible: leading formulation, evidence for it, strongest alternative, and the result that would separate them.

The habit worth carrying out of this chapter is a **longitudinal formulation** rather than a fixed label. A hormone level moves with treatment, a deficiency corrects and recurs, a metabolic burden accumulates over years of prescribing, and a syndrome once fully explained by a systemic disease may later be independent of it. Keep revising in the open: retain what remains supported, discard what no longer fits, and be explicit with the patient about which part of the picture is now being treated and why.

> Establish the direction of causation, name the organ from the somatic signature rather than the mental state, ask whether a treatment window is open, test the question rather than the panel, and hand over constructs owned by another chapter.

2 · HPA axis physiology and its role in psychiatric disorders

This is the axis every other endocrine question in psychiatry hangs from. Examiners return to it because it asks for three kinds of knowing at once: physiology recited accurately, laboratory findings among the most replicated in the discipline, and the interpretive discipline that stops those findings being over-read. Get the physiology right and much of what follows in adrenal disease, steroid psychiatry and the biomarker literature becomes derivable rather than memorised. A candidate who cannot say which adrenal output answers to the pituitary cannot explain why one kind of adrenal failure needs mineralocorticoid replacement and another does not.

Begin with what the gland makes, because this axis is only one of several control systems running through it. Kaplan and Sadock's Comprehensive Textbook of Psychiatry frames the adrenal as two organs in one capsule: steroids from the cortex and catecholamines from the medulla. Within the cortex, three hormone classes answer to different masters. Glucocorticoids, principally cortisol, are driven by ACTH from the anterior pituitary under a negative feedback loop; aldosterone, a mineralocorticoid, also arises from the cortex but is controlled by the renin-angiotensin system, itself influenced by body volume and potassium balance. The adrenal androgens, predominantly **dehydroepiandrosterone**, are regulated by the ACTH system and undergo peripheral conversion to sex-determining androgens. Only the first is the HPA axis proper, and keeping the three apart makes the rest of this chapter coherent.

8 AM	ABOUT 15 MIN	20 TO 40%	40 TO 60%
PEAK OF CIRCULATING GLUCOCORTICOID	PLASMA HALF-LIFE OF ACTH	DEPRESSED OUTPATIENTS, RAISED HPA ACTIVITY	DEPRESSED INPATIENTS, RAISED HPA ACTIVITY

2.1 Three modes of regulation

Goodman and Gilman's The Pharmacological Basis of Therapeutics compresses the whole of ACTH regulation into three characteristic modes. Learn them as three, because almost every abnormality you will meet is a failure of one of them.

- **A diurnal rhythm in basal steroidogenesis.** The gland runs to a clock, which is what makes any single measurement meaningless without its timestamp.
- Negative-feedback regulation by adrenal corticosteroids. The product restrains its own drivers, turning an open cascade into a governed system.
- Marked increases in steroidogenesis in response to stress. A designed feature, not a disturbance, and the mode that most often contaminates a clinical measurement.

The framing earns its keep by separating questions that look alike. Whether cortisol is high, whether the rhythm is intact and whether feedback works are three questions needing different measurements under different conditions, and they dissociate. The accompanying figure sets the loop out in the order the modes act on it.

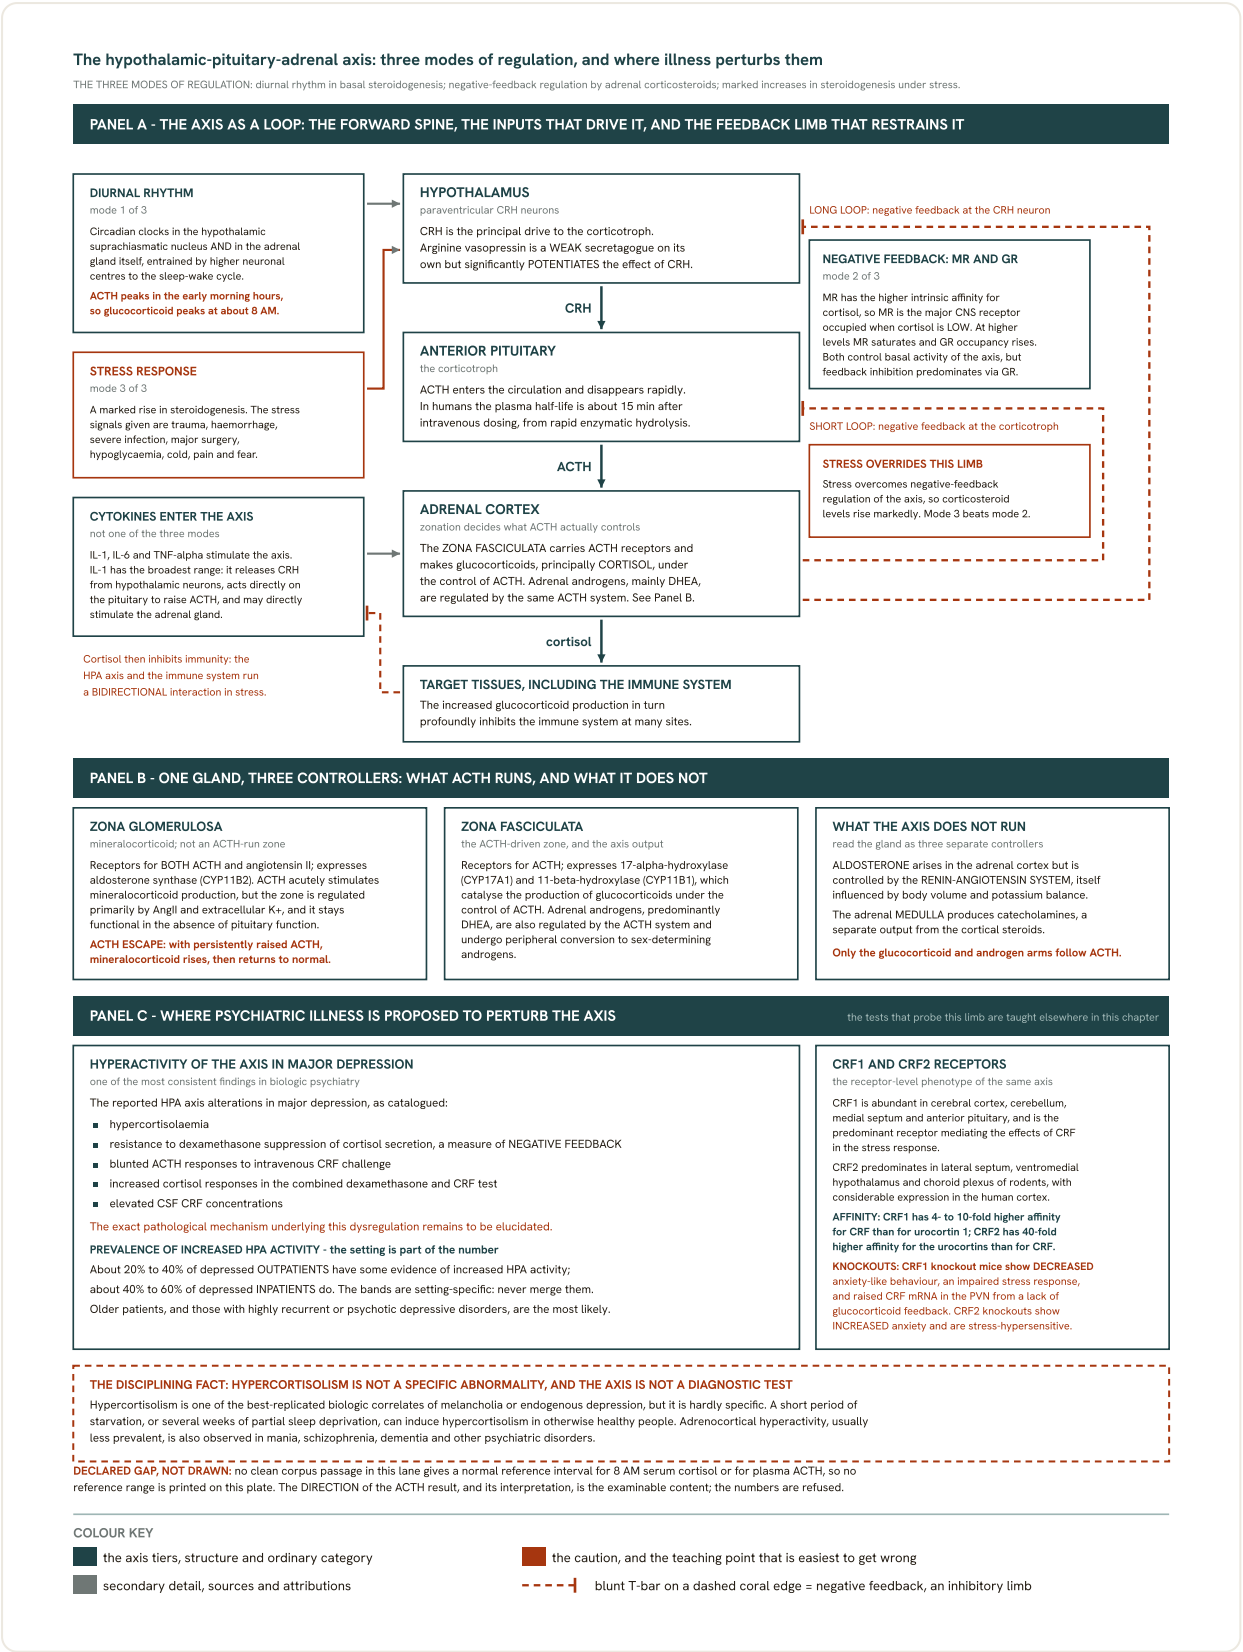


Fig 33.1 The hypothalamic-pituitary-adrenal axis as a loop: the forward spine, the feedback limb, and where illness is proposed to perturb it. Three modes govern the axis: a diurnal rhythm in basal steroidogenesis, negative feedback by adrenal corticosteroids, and a marked rise under stress. Read the feedback limb as the examinable structure: predominantly glucocorticoid-receptor mediated, acting at both the corticotroph and the CRH neuron, and overridden by stress. Hyperactivity of the axis in major depression is consistent but not specific. (Goodman and Gilman's The Pharmacological Basis of Therapeutics; Kaplan and Sadock's Comprehensive Textbook of Psychiatry 2024.)

MNEMONIC

Rhythm, Feedback, Stress

The three modes of physiological regulation, in the order they are usually examined. **R**hythm gives the timing rule, **F**eedback gives the receptor story and the logic of suppression testing, and **S**tress gives the reason a raised value so often means nothing at all.

2.2 The diurnal rhythm and the morning anchor

The rhythm is not generated at one site. It is set by circadian clocks in the hypothalamic **suprachiasmatic nucleus** and in the adrenal gland itself, and the pair is entrained by higher neuronal centres responding to the sleep-wake cycle. Because the adrenal carries its own oscillator, the rhythm is not merely pituitary drive relayed downward; because entrainment runs through sleep and waking, it belongs to the patient's schedule rather than to the wall clock.

The consequence a clinician uses daily is timing. ACTH peaks in the early morning hours and circulating glucocorticoid follows, peaking at about **8 AM**. Every convention about when to draw a cortisol sample descends from that one fact, as does the logic of the suppression testing this chapter takes up where it deals with Cushing syndrome and the HPA biomarkers. Asked why a morning sample is specified, the answer is not tradition: only a measurement at the reproducible peak of a rhythmic variable is comparable between patients or visits.

GOOD TO REMEMBER

Before interpreting any cortisol value, ask what the patient's night looked like. Because the rhythm is entrained by the sleep-wake cycle, a patient who worked a night shift, travelled overnight or slept badly on a ward is not on the schedule the convention assumes. The sample is timed to the clock; the physiology is timed to the patient.

2.3 Negative feedback: a receptor handover, not a thermostat

Cortisol acts centrally through two receptor species, and the logic of feedback rests on the difference in their affinities. The **mineralocorticoid receptor** has the higher intrinsic affinity for cortisol, so during periods of low cortisol it is the major receptor species occupied in the central nervous system. As blood cortisol rises that receptor saturates and occupancy of the **glucocorticoid receptor** increases. Both control basal activity of the axis, but feedback inhibition by glucocorticoids predominantly occurs through the glucocorticoid receptor.

Read it as a handover rather than a switch. At low levels the system runs on the high-affinity receptor and is tuned for maintenance; as levels climb the low-affinity receptor is recruited and the braking function comes on with it. That is why a graded system behaves differently at different set points, and it is the mechanism the Maudsley prescribing guidance uses to explain the neuropsychiatric effects of exogenous steroids, taken up in this chapter under steroid psychiatry. A drug that occupies one receptor species preferentially does not simply raise a level; it changes the balance of the two signals the brain is reading.

RULE **Basal tone is set by both receptors; braking is predominantly a glucocorticoid-receptor job.** The mineralocorticoid receptor is occupied first because its affinity for cortisol is higher, and the glucocorticoid receptor is recruited only as levels rise and the first saturates. Any account of impaired feedback in disease must say which limb of that handover is at fault.

2.4 Zonation: why the two halves of the cortex fail differently

The cortex is not uniformly ACTH-dependent, the most examinable piece of adrenal anatomy. Cells of the **zona glomerulosa** carry receptors for both ACTH and angiotensin II and express aldosterone synthase, the enzyme CYP11B2; ACTH acutely stimulates mineralocorticoid production there, but the zone is regulated primarily by angiotensin II and extracellular potassium and remains functional in the absence of pituitary function. Cells of the **zona fasciculata** carry ACTH receptors and express 17-alpha-hydroxylase (CYP17A1) and 11-beta-hydroxylase (CYP11B1), which catalyse glucocorticoid production under the control of ACTH.

That asymmetry explains a clinical pattern this chapter develops where it deals with adrenal insufficiency: failure at pituitary level removes the ACTH-dependent glucocorticoid limb while leaving largely intact a zone that never needed the pituitary, whereas destruction of the gland removes both. Hold the anatomy and the consequence need not be memorised.

AXIS OF COMPARISON	ZONA GLOMERULOSA	ZONA FASCICULATA
Receptors expressed	ACTH and angiotensin II	ACTH
Characteristic enzymes	Aldosterone synthase (CYP11B2)	CYP17A1 and CYP11B1
Principal product	Mineralocorticoid	Glucocorticoid
Chief regulator	Angiotensin II and extracellular potassium	ACTH
Behaviour without a pituitary	Remains functional	Loses its drive

One further phenomenon is regularly misremembered. With persistently elevated ACTH, mineralocorticoid levels initially increase and then return to normal, a behaviour named **ACTH escape**: the zone responds acutely and then reverts to the control system that actually governs it.

TRAP **Reasoning that chronically high ACTH must produce sustained mineralocorticoid excess.** The premise is right and the conclusion wrong. Candidates correctly recall that the outer zone bears ACTH receptors and that ACTH acutely stimulates mineralocorticoid output, then extrapolate to a chronic effect. Escape is the answer: the initial rise is not maintained, because angiotensin II and extracellular potassium, not the pituitary, are the standing regulators of that zone.

2.5 Driving the axis upward: vasopressin and the stress override

Corticotrophin-releasing hormone is the principal drive to the corticotroph, but it does not work alone. **Arginine vasopressin** is a weak secretagogue for corticotrophs on its own, yet it significantly potentiates the effects of corticotrophin-releasing hormone. It is a modulator rather than an independent driver, so conditions that raise it amplify an axis already being driven.

The larger point about upward drive is that the system is built to be overridden. Stress overcomes negative feedback regulation of the axis and produces a marked rise in corticosteroid levels, and the qualifying signals are mostly not psychological in the everyday sense: trauma, haemorrhage, severe infection, major surgery, hypoglycaemia, cold, pain and fear. A hypoglycaemic, febrile, post-operative or frightened patient has a legitimately elevated cortisol, with nothing in the axis malfunctioning at all.

PITFALL

Sampling the axis in a patient who is acutely unwell, then attributing the result to a primary endocrine or psychiatric process. A raised cortisol drawn during sepsis, uncontrolled pain, hypoglycaemia or the first hours after surgery measures the stress response, not the set point: the physiology says the value should be high. Repeat it when the patient is no longer a stressed patient, or do not draw it at all.

2.6 Corticotrophin-releasing factor receptors and the behavioural phenotype

Receptor pharmacology is where this axis stops being purely endocrine and becomes a plausible substrate for anxiety and mood. Kaplan and Sadock's Comprehensive Textbook of Psychiatry describes two subtypes with different distributions, different ligand preferences and opposite behavioural phenotypes, a contrast favoured by written papers because it is clean enough to test.

AXIS OF COMPARISON	CRF1 RECEPTOR	CRF2 RECEPTOR
Principal distribution	**Cerebral cortex, cerebellum, medial septum and anterior pituitary**	**Lateral septum, ventromedial hypothalamus and choroid plexus in rodents, with considerable expression in human cortex**
Ligand preference	4- to 10-fold higher affinity for corticotrophin-releasing factor than for urocortin 1	40-fold higher affinity for the urocortins relative to corticotrophin-releasing factor
Role in the stress response	Appears to be the predominant receptor mediating the effects of corticotrophin-releasing factor	Not the predominant mediator
Knockout phenotype	Decreased anxiety-like behaviour, impaired stress response, elevated corticotrophin-releasing factor messenger RNA in the paraventricular nucleus from absent glucocorticoid negative feedback	Increased anxiety-like behaviour and hypersensitivity to stress

Two cautions travel with the table. The knockout phenotypes are animal data, and the species caveat is in the source itself: the rodent pattern for the second subtype does not transfer wholesale to humans, where cortical expression is substantial. Nor are the opposing phenotypes a push and pull on one dial, since the elevated messenger RNA in the first knockout is a feedback

phenomenon, a hypothalamic response to missing glucocorticoid restraint rather than a direct effect of losing the receptor.

2.7 The immune interface

Now the part that turns a hormone axis into a systems story. The relationship with the immune system runs both ways, and holding only one direction is the commonest incompleteness in candidates' answers. Among the proinflammatory cytokines, interleukin-1, interleukin-6 and tumour necrosis factor alpha stimulate the axis, with interleukin-1 having the broadest range of actions, and that breadth means action at all three levels at once: interleukin-1 stimulates release of corticotrophin-releasing hormone by hypothalamic neurons, acts directly at the pituitary to increase ACTH release, and may directly stimulate the adrenal to produce glucocorticoids. The return arm closes the loop, since the resulting increase in glucocorticoid production profoundly inhibits the immune system at multiple sites. Goodman and Gilman's The Pharmacological Basis of Therapeutics states the conclusion plainly: axis and immune system are capable of bidirectional interactions in response to stress.

Two consequences matter for a psychiatrist. An inflammatory illness is an endocrine event as well as an infective one, a further reason the acutely septic patient is a poor subject for axis measurement. And this is the mechanistic bridge inflammatory hypotheses of depression are built on, since a route from peripheral immune activation to the hypothalamus is what such a hypothesis requires. That the bridge exists does not establish that it carries the traffic.

2.8 The axis in psychiatric disorder, and the limits of the finding

Hyperactivity of the axis in major depression remains one of the most consistent findings in biological psychiatry. Be precise about what has been reported, because it is a family of observations rather than one. The catalogue comprises hypercortisolemia; resistance to dexamethasone suppression of cortisol secretion, which is a measure of negative feedback; blunted ACTH responses to intravenous corticotrophin-releasing factor challenge; increased cortisol responses in the combined dexamethasone and corticotrophin-releasing factor test; and elevated cerebrospinal fluid corticotrophin-releasing factor concentrations. The tests are set out in this chapter where it deals with the HPA biomarkers; what belongs here is that each probes a different arm of the physiology above, and that the mechanisms underlying the dysregulation remain to be elucidated.

How common is it? The honest answer is setting-dependent, and the two figures must never be collapsed into one. About **20% to 40%** of depressed outpatients show some evidence of increased activity of the axis, against about **40% to 60%** of depressed inpatients. Older patients, particularly those with highly recurrent or psychotic depressive disorders, are the most likely to show it. A single pooled percentage quoted without its setting is not a shorter version of those two numbers but a different and unsupported claim, and the gradient is itself informative: the inpatient population is both more severely ill and more likely to be under the acute stress described above.

The disciplining fact governs everything the biomarker literature may claim. Although hypercortisolism is one of the best-replicated biological correlates of melancholia or endogenous depression, it is hardly a specific abnormality. Two everyday states produce it in people who are entirely well: a short period of starvation, or several weeks of partial sleep deprivation. Nor is it confined to depression among psychiatric populations, since adrenocortical hyperactivity, although usually less prevalent, is observed in mania, schizophrenia, dementia and other psychiatric disorders. Consistency and specificity are different properties, and this finding has a great deal of the first and very little of the second.

TRAP **Treating hypercortisolism as a diagnostic test for depression.** The reasoning is superficially sound: the finding is among the most replicated in biological psychiatry, therefore it must be diagnostically useful. Replication establishes that an association is real, not that it discriminates. A fasting volunteer and a chronically sleep-deprived resident both produce the abnormality, and it appears across mania, schizophrenia and dementia. State the finding, then state its non-specificity in the same breath; an answer that gives only the first half is the answer being tested for.

2.9 Interrogating the axis: what the physiology dictates before any test

Management of the adrenal disorders and of steroid-induced syndromes is set out in this chapter where each condition is taken up; what the axis itself dictates is the assessment logic.

LOCUS decides what is realistically obtainable, more sharply here than for most investigations. In an outpatient clinic a patient can be brought back unhurried and on their usual sleep schedule for a sample timed to the morning peak, which is the condition under which the number means what the convention assumes. On an acute ward sleep is disrupted and the illness that prompted admission is often itself a stress signal, so the convention is met on paper while the physiology it rests on is not. In an emergency setting it should not be interrogated for psychiatric purposes at all, since the setting guarantees a stress response.

FOCUS changes with phase. Before any blood is drawn the target is the timestamp and the state of the patient. Acutely, the only defensible question is whether an adrenal emergency is in play, a medical rather than a psychiatric determination. Over the short and medium term the questions are the ones this chapter's later sections are built on: is the level abnormal, is the rhythm preserved, does feedback still work. Long term, in a patient on exogenous steroid, the question is what sustained occupancy of the two receptor species is doing to mood and cognition.

MODUS is the set of routes available, each mapping onto one regulatory mode. A timed morning sample addresses the basal level at a reproducible point in the rhythm; repeat sampling across the waking day addresses the rhythm itself; pharmacological suppression addresses feedback, and dynamic challenge with corticotrophin-releasing factor addresses pituitary responsiveness. Both of the last two are taught where this chapter deals with the HPA biomarkers. One pre-analytical constraint belongs to the physiology rather than to any protocol: ACTH is readily absorbed from parenteral sites and disappears rapidly from the circulation, with a plasma half-life in humans of about **15 minutes**, primarily from rapid enzymatic hydrolysis. Hence the need to handle an ACTH sample promptly.

IN INDIA

Two divergences bite in routine practice. Cortisol assays are commonly run in-house while plasma ACTH is frequently a send-out test, so the sample is drawn where the patient is and assayed wherever the analyser sits. The same rapid enzymatic hydrolysis that gives ACTH its short plasma half-life is why prompt handling and transport are not optional refinements: a sample that spent the afternoon on a bench yields an uninterpretable number that will nonetheless be acted upon, and where a courier interval cannot be guaranteed it is better to arrange the test through a laboratory that can. Second, reference intervals for these hormones are assay-dependent and are not reproduced in this chapter. Read every result against the interval printed by the performing laboratory, and expect values from two laboratories not to be comparable.

Three modes of regulation, one timing rule and one caveat: rhythm, feedback and the stress override; sample at the morning peak because that is where the rhythm is reproducible; and remember that raised cortisol is consistent in depression but specific to nothing.

3 · Cushing syndrome, steroid psychiatry and the HPA biomarkers

Two clinical realities sit under one heading, and only one is rare. Endogenous hypercortisolism a psychiatrist meets a few times in a career; corticosteroid-induced disturbance is on the wards every week, because exogenous steroid, usually prescribed for a steroid-sensitive medical condition, is the commonest cause of hyperadrenalism overall.

UP TO 70%	**OVER 50%**	**22.2**	**11.5 DAYS**
DEPRESSION IN CUSHING SYNDROME, CONTESTED	COGNITIVE IMPAIRMENT IN CUSHING SYNDROME	NEUROPSYCHIATRIC CASES PER 100 PERSON-YEARS, FIRST COURSE	MEDIAN TIME TO ONSET ON CORTICOSTEROIDS

3.1 Where the excess cortisol comes from

Hyperadrenalism is excess adrenocortical corticosteroid. Usually that is ACTH overproduction by the anterior pituitary, most often from a pituitary adenoma, the subset called **Cushing disease**; the remainder is ACTH secretion by a non-endocrine tumour, a primary adrenal neoplasm, and administered steroid.

The somatic picture: hypertension, muscle weakness and fatigue, osteoporosis, cutaneous striae and easy bruising, and obesity of the upper face, back and mesentery giving moon facies, buffalo hump and truncal obesity. Biochemically it declares itself twice: plasma and urinary cortisol are elevated and the normal circadian rhythm of secretion is blunted or absent. Lost rhythm, not a high level alone, is what the tests exploit; the feedback anatomy is taught where this chapter covers the physiology of the hypothalamic-pituitary-adrenal axis.

3.2 The psychiatric signature of endogenous hypercortisolism

The dominant picture is depressive, not psychotic. Depression is the most frequently noted psychiatric symptom, in up to 70% of patients, and may precede the physical features: the patient reaches psychiatry first. Nor is it merely comorbid: it responds to treatment of the underlying cause, the response correlating with the fall in plasma cortisol.

Anxiety is commonly comorbid in up to 50% of those who are depressed, the denominator being the depressed subgroup, not every patient. Elevated mood is infrequent, under 10% in most case series, and delirium is also uncommon, usually marking a supervening infection or another metabolic disturbance such as metabolic alkalosis. New confusion is a reason to hunt a second illness, not to blame the first harder. Delirium as a syndrome belongs to the Stroke, TBI, delirium and focal brain syndromes notes.

Psychosis is where folk teaching misleads. Reports of Cushing syndrome presenting with psychosis and obscuring the underlying condition exist, but that presentation is rare, and psychotic symptoms are almost always **mood-congruent**: delusional beliefs and derogatory auditory hallucinations with depressed mood. Cognition is affected far more often and examined far less: impairment occurs in over 50%, in verbal memory, attention, and visuomotor and visuospatial function.

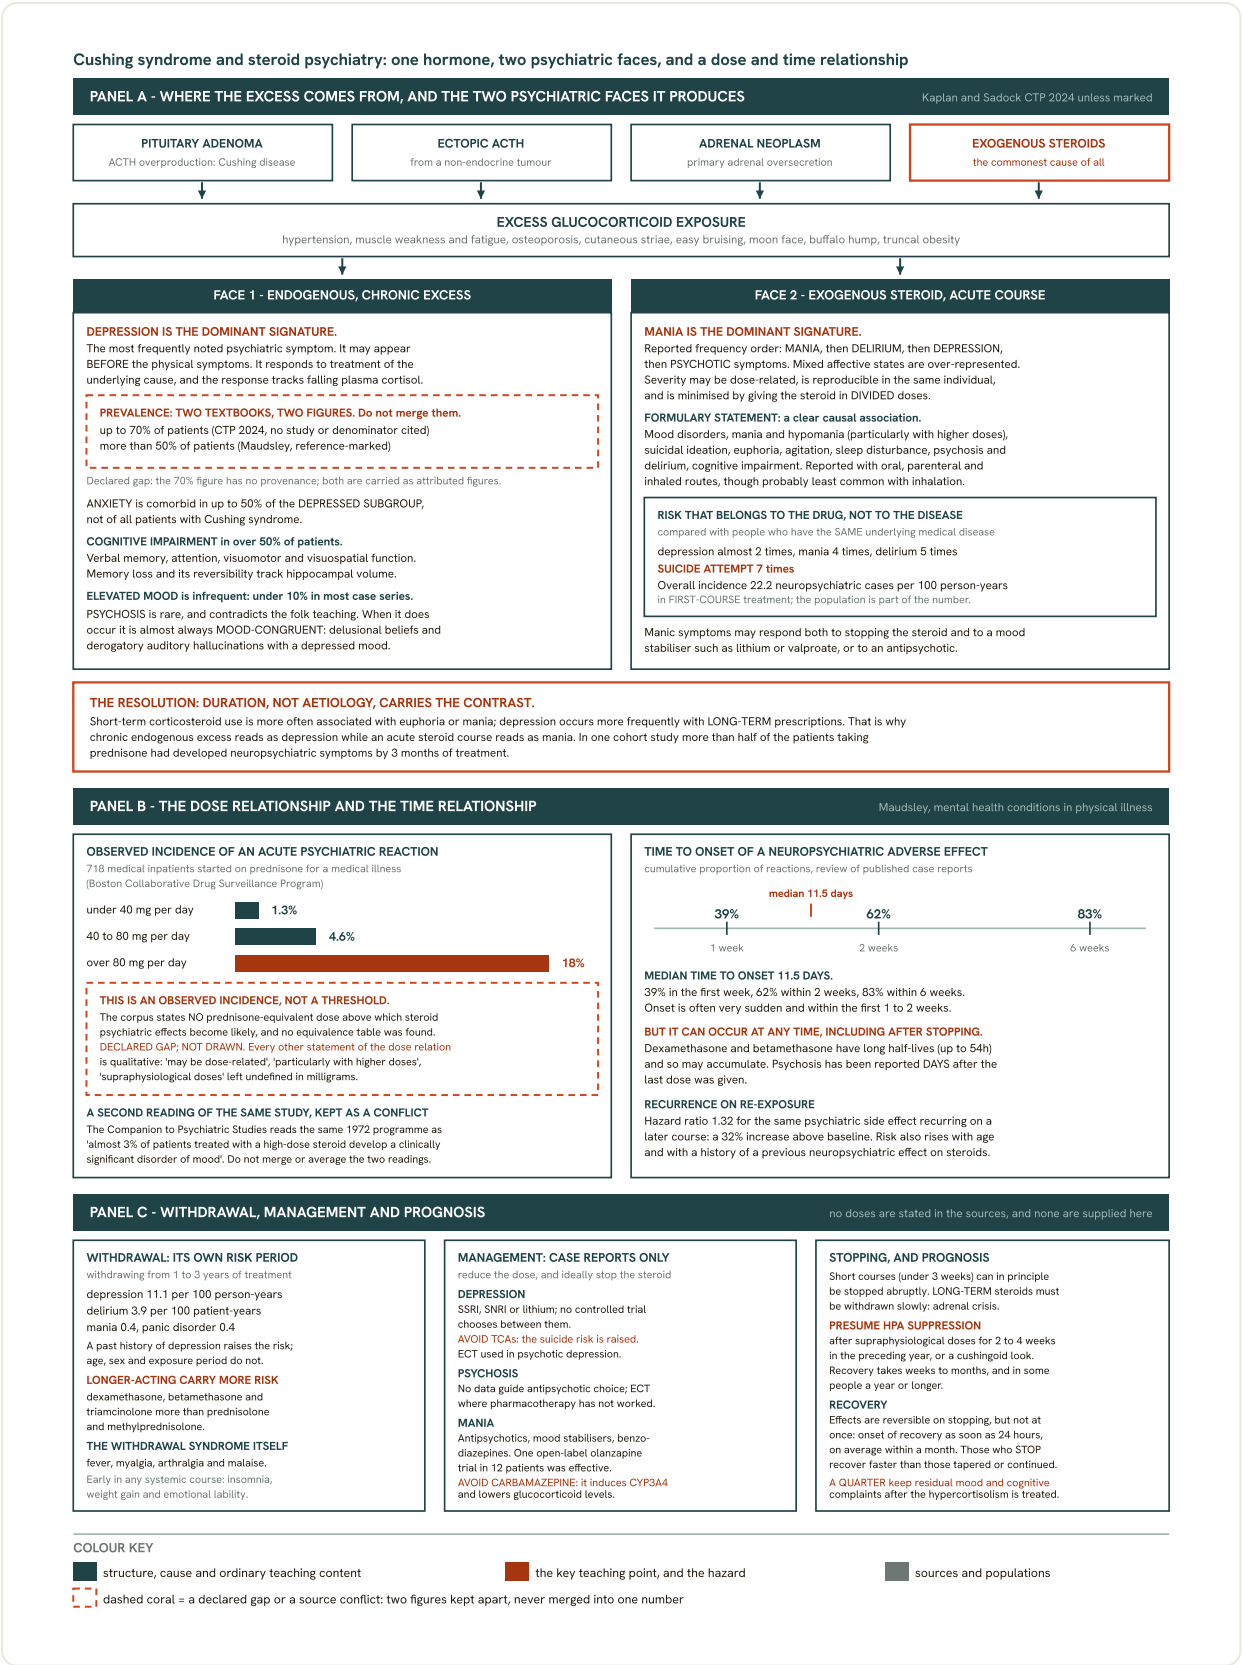


Fig 33.2 **Cushing syndrome and steroid psychiatry: two faces of one hormone.** Chronic endogenous excess reads as depression with cognitive impairment; an acute exogenous course reads as mania, delirium and mixed states, and duration rather than aetiology carries the contrast. The dose bands are observed incidences in inpatients started on prednisone for a medical illness, not a clinical threshold: no prednisone-equivalent threshold exists in the corpus, and that gap is marked rather than filled. (Kaplan and Sadock CTP 2024; Maudsley; Companion to Psychiatric Studies; Goodman and Gilman.)

The two prevalences do not agree. The Comprehensive Textbook of Psychiatry gives up to 70% with no study, denominator or citation; the Maudsley Prescribing Guidelines for Mental Health Conditions in Physical Illness states instead that depression occurs in more than 50% of patients with Cushing's disease, and does carry a reference. The ranges overlap and the populations differ, so neither is safe as the prevalence: cite each with its textbook and population, or say "the majority".

3.3 Prescribed steroids: a different syndrome, or a different duration?

Swap the disease for a prescribed drug and the affective direction reverses. With exogenous steroid, manic symptoms are most frequently reported, followed by delirium, depression and psychotic symptoms, and mixed affective states appear over-represented. The tidy conclusion, that endogenous cortisol depresses and synthetic steroid elevates, is wrong: the variable is duration. Short-term use is more often associated with euphoria or mania, depression more frequently with long-term prescriptions, consistent with depression being common in Cushing's disease, a disorder of chronic endogenous excess; synthetic corticosteroids have the same effect, and in one cohort study more than half of patients on prednisone had neuropsychiatric symptoms by 3 months.

RULE **Duration, not aetiology, carries the direction of mood change in glucocorticoid excess.** A long course travels from the elevating end of that curve to the depressing end without changing drug; endogenous Cushing syndrome is the chronic arm.

MNEMONIC

Short lifts, long sinks

No threshold in weeks is licensed, because none is stated.

The formulary adds route and severity: recognised glucocorticoid psychiatric adverse effects are mood disorders, mania and hypomania particularly at higher doses, suicidal ideation, euphoria, agitation, sleep disturbance, psychosis, delirium, dementia and cognitive impairment; the causal association is clear and most substantial for depression and mania; they are reported with oral, parenteral and inhaled routes, though probably less common with inhalation, and generally respond to dose decreases. Inhaled steroid is less implicated, not exempt, so respiratory prescriptions belong in the drug history.

3.4 Dose, incidence, and one study read two ways

Risk rises with dose, and the only numeric stratification comes from one surveillance programme. Among 718 medical inpatients started on prednisone for medical illness, acute psychiatric reactions occurred in 1.3% of those on less than 40 mg prednisone per day, 4.6% at 40 to 80 mg per day, and **18%** at more than 80 mg per day. Those are incidences of adverse reaction within dose bands, not therapeutic doses.

TRAP **Two standard textbooks summarise the same surveillance programme with figures differing roughly six-fold, and both circulate in examination answers.** The Maudsley guidelines give the dose-stratified incidences above; the Companion to Psychiatric Studies renders the same programme as corticosteroids increasing the risk of affective disturbance, and of mania or hypomania in particular, in relation to dose, with almost 3% of patients on a high-dose steroid developing a clinically significant disorder of mood. Both cannot describe the high-dose stratum: either the second is an overall cohort rate misattached to that group, or the two count different outcomes, acute psychiatric reaction against clinically significant mood disorder. Do not pick one and do not average them: cite the dose-stratified version with its source and note the second text's almost 3% alongside it.

Two caveats. No source states a prednisone-equivalent threshold above which psychiatric effects become likely; the most direct statement is that severity may be dose-related, reactions are reproducible in the individual, and they may be minimised by divided doses. An observed incidence within a range is not a threshold. Second, elevation is commoner than illness: controlled studies confirm high rates of subclinical mood elevation after 1 week of steroid administration.

Population data give the sturdier framing. A large United Kingdom database study put the incidence of neuropsychiatric outcomes at 22.2 cases per 100 person-years for first-course treatments, its internal comparison attributing the effect to the drug rather than the illness: against people with the same underlying disease, glucocorticoid-treated patients were almost twice as likely to become depressed, four times as likely to develop mania, five times delirium and seven times to attempt suicide. Prior experience counts too: risk rises with age and with previous neuropsychiatric adverse effects on corticosteroids, and one study put the hazard ratio for recurrence of the same side effect on a later course at 1.32, a 32% increase above baseline.

3.5 Timing: onset, the tail, and withdrawal

Timing is what makes this diagnosis on a ward round. Effects can appear at any point, from almost immediately after initiation to some time after treatment has stopped, but most fall in the first few weeks: a case-report review found a median onset of **11.5 days**, with 39% in the first week, 62% within 2 weeks and 83% within 6 weeks. The formulary agrees that onset is often very sudden and within the first 1 to 2 weeks. A new manic or psychotic presentation in a medical inpatient started on steroid a fortnight ago is steroid-related until proven otherwise.

The tail runs longer than intuition allows. Dexamethasone and betamethasone have long half-lives, up to 54 hours, and may accumulate, and psychosis has been reported days after the last dose. A drug chart showing the steroid already stopped does not exonerate it. Withdrawal is its own risk window: in a United Kingdom primary care study of withdrawal from 1 to 3 years of treatment, incident depression ran at 11.1 per 100 person-years, ahead of delirium at 3.9 per 100 patient-years, mania at 0.4 and panic disorder at 0.4; previous depression raised the risk while age, sex and exposure period did not, and the longer-acting dexamethasone, betamethasone and triamcinolone carried more risk than shorter-acting prednisolone and methylprednisolone.

Across a whole course: early therapy brings insomnia, weight gain and emotional lability; high-dose or long-term therapy brings psychosis alongside infection risk, osteoporosis, osteonecrosis, myopathy and suppression of the axis itself; and cessation brings acute hypocortisolism.

3.6 Mechanism, and what does not fully reverse

The honest position is that the mechanism is not clear. The best-articulated account is receptor balance: exogenous corticosteroid activates pituitary glucocorticoid receptors while suppressing adrenal cortisol secretion, depleting endogenous cortisol and reducing mineralocorticoid receptor stimulation, and that imbalance between glucocorticoid and mineralocorticoid receptor activation may underlie the neuropsychiatric effects; high cortisol also inhibits brain-derived neurotrophic factor, low levels of which may contribute to depression and anxiety. The receptor affinities behind the imbalance are taught where this chapter covers the physiology of the axis.

The cognitive picture has an anatomical correlate. The degree of memory impairment and its reversibility correlate with hippocampal volume, implicating the effect of excess glucocorticoid on hippocampal neurons, and a more general reduction in grey matter volume, cortical thickness and white matter integrity can persist after the hypercortisolism is corrected. Imaging converges: smaller hippocampal, cingulate gyrus and medial frontal gyrus volumes, with computed tomography, magnetic resonance imaging and autopsy showing global cerebellar and cerebral volume loss in both endogenously and exogenously exposed patients; a quarter keep residual mood and cognitive complaints after treatment, and brain volume often does not return completely.

> **GOOD TO REMEMBER**
>
> Correcting the cortisol is not the same as restoring the patient. Say so at the outset rather than letting a patient discover it months after curative surgery, and record the cognitive deficits in your own descriptive terms.

3.7 Managing corticosteroid-induced psychiatric syndromes

Guidance comes principally from the Maudsley guidelines, with Goodman and Gilman's The Pharmacological Basis of Therapeutics governing withdrawal and adrenal safety. Neither states a psychotropic dose here, and none should be imported.

LOCUS. Mostly liaison work on a medical ward, because the steroid is treating something that needs treating. Emergency presentation is delirium, acute mania and self-harm risk, real given the seven-fold excess of suicide attempts above. Outpatient work is surveillance and safe tapering on maintenance. The psychiatrist rarely controls the steroid prescription and must negotiate it.

FOCUS. Acutely, agitation, psychotic symptoms, manic elevation and immediate safety. Short term, the steroid dose itself. Medium term, recovery and re-exposure, weighted by the recurrence hazard. Long term, safe withdrawal, adrenal recovery and residual symptoms.

MODUS. The first modality is the offending drug. Dose reduction, and ideally cessation, is the preferable strategy when psychiatric side effects occur; short courses of less than 3 weeks can in principle be stopped abruptly, whereas long-term corticosteroids must be withdrawn slowly to

minimise the risk of adrenal insufficiency and crisis. Pharmacotherapy is evidence-poor and largely a matter of what to avoid. In corticosteroid-induced depression, case reports support SSRIs, SNRIs and lithium with no controlled trials to direct the choice, tricyclics are avoided given the heightened suicide risk, and electroconvulsive therapy has been used in psychotic depression. In psychosis no data guide antipsychotic selection, tricyclics are again avoided since some case series suggest they worsen these psychoses, and electroconvulsive therapy has succeeded where drugs failed. In mania, case reports describe antipsychotics, mood stabilisers, benzodiazepines and combinations, one open-label trial of olanzapine in 12 patients with mania or mixed symptoms found it effective and well tolerated, and carbamazepine is avoided because CYP3A4 induction lowers plasma glucocorticoid concentrations. Lithium as a drug belongs to the Mood disorders: pharmacotherapy notes and antipsychotic selection to the Schizophrenia: management, antipsychotics and adverse effects notes.

Prognosis is good, with a caveat about speed. These effects are reversible on stopping treatment but do not necessarily resolve immediately: recovery from corticosteroid-induced psychosis could begin as soon as 24 hours after either treatment for the psychosis was started or the steroid was stopped, occurred on average within a month, and was quicker in patients able to discontinue than in those tapered or continued. That comparison is the argument to take to the referring team, and the one adrenal risk constrains.

PITFALL

The reflex on seeing steroid-induced psychosis is to stop the steroid today. On a long course that can precipitate adrenal crisis. Patients given supraphysiological glucocorticoid doses for 2 to 4 weeks within the preceding year, or who look cushingoid, should be presumed to have some axis impairment until proven otherwise; acute adrenal insufficiency follows overly rapid withdrawal after prolonged therapy has suppressed the axis, and recovery takes several weeks to months in many and a year or longer in some. Recognise the **glucocorticoid withdrawal syndrome** of fever, myalgia, arthralgia and malaise rather than reading it as somatisation. No source defines supraphysiological in milligrams.

3.8 The dexamethasone suppression test: one name, several tests

Dexamethasone probes negative feedback: a potent synthetic glucocorticoid at night, and a normal axis shuts cortisol output down by morning. Failure to do so is **nonsuppression**. That one idea has been implemented as at least four procedures, and the differences are the examination.

The psychiatric version made the test famous: to confirm a diagnostic impression of major depression, the patient takes 1 mg of dexamethasone by mouth at 11 PM with plasma cortisol at 8 AM, 4 PM and 11 PM, a value above **5 microgram/dL** counting as nonsuppression and therefore positive. The endocrine screen shares the protocol but not the threshold. **The overnight dexamethasone suppression test** is the screening procedure of choice for Cushing syndrome: 1 mg of dexamethasone at midnight with a morning plasma cortisol, a level greater than **200 nmol per litre** indicating a high likelihood of the syndrome; it outperforms urinary

free cortisol on sensitivity and specificity, whereas an isolated plasma cortisol screens inadequately because levels vary considerably across the day.

THE QUESTION BEING ASKED	DOSE TIMING	CORTISOL SAMPLING	THRESHOLD CALLED ABNORMAL	SOURCE
Confirming a diagnostic impression of major depression	11 PM	8 AM, 4 PM and 11 PM	Above 5 microgram/dL	Comprehensive Textbook, glossary
Screening for Cushing syndrome	Midnight	Following morning	Greater than 200 nmol per litre	Comprehensive Textbook, endocrine text
Biochemical evidence of raised cortisol biosynthesis	11 PM	8 AM following morning	Markedly lower than either above; not recoverable from the source	Goodman and Gilman
Research probe of feedback inhibition in depression	Suppression normally lasts 24 hours	8:00 AM, escape at 4:00 PM or 11:00 PM	Nonsuppression, or later escape	Comprehensive Textbook, neuroscience text

TRAP **Quoting a suppression-test cut-off without naming the question the test was asked.** Three textbook passages, two from the same textbook, give three post-dexamethasone thresholds for a test all three call the overnight 1 mg test, and two dosing times; they are not one number in different units. Likeliest is that the psychiatric test of the 1980s and the modern endocrine screen are different tests sharing a protocol, the endocrine threshold set low for sensitivity, but that is a hypothesis and not a resolution: do not harmonise them. The dose is contested the same way: a research probe of feedback inhibition in depression uses dexamethasone at 0.5 to 2.0 mg, whereas the diagnostic protocol in suspected hypercortisolism uses 1 mg of dexamethasone orally at 11 PM with cortisol at 8 AM the following morning. Say 1 mg is the standard clinical protocol and the research literature has used a range; dropping either loses the mark.

A fourth procedure answers a different question. **The high-dose dexamethasone suppression test** determines the aetiology of biochemically documented Cushing syndrome: after 48 hours of baseline cortisol, dexamethasone is given orally as 2 mg every 6 hours for 48 hours, or 8 mg overnight, and in many though not all patients with a pituitary source of ACTH excess cortisol will suppress. The gold standard for establishing a pituitary source remains **inferior petrosal sinus sampling** after CRH administration. This localises within confirmed disease, and is never a screening test nor a psychiatric one.

3.9 The axis as a biomarker, and the other probes

The verdict should anchor everything above: neither the sensitivity nor the specificity of these tests of feedback inhibition is sufficient for use as a diagnostic test, and adrenocortical hyperactivity, although usually less prevalent, is seen in mania, schizophrenia, dementia and other psychiatric disorders. The source gives no sensitivity or specificity figures and none should

be invented. The failure is chiefly one of specificity, since nonsuppression also occurs in major depression, alcoholism, pregnancy, and during acute physical or emotional stress.

- Medication: phenytoin, barbiturates, carbamazepine.
- Endocrine factors: Cushing's disease or pregnancy.
- Metabolic problems: recent withdrawal from alcohol, rapid weight loss, malnutrition, or nausea and vomiting.
- Multiinfarct dementia and increased intracranial pressure.

The enzyme inducers are mechanistically coherent, since drugs or conditions accelerating dexamethasone metabolism, and drugs such as estrogens or conditions such as pregnancy that raise cortisol-binding globulin, compromise the test. Vascular cognitive disorder appears there as a false-positive cause only; the dementia syndrome belongs to the Dementias and neurocognitive disorders notes. Beneath the interpretive problem lies a measurement one: cortisol hypersecretion and dexamethasone nonsuppression correlate only imperfectly, at approximately 60% concordance, so they measure overlapping but different aspects of axis dysfunction.

Where the axis earns its keep is prognostic. Hypercortisolism usually resolves with effective treatment, and when it persists despite it, increased axis activity carries a high risk of relapse, reflecting incomplete resolution of the depressive pathophysiology; such patients respond less well to attention placebo interventions and psychotherapy and may need active pharmacotherapy or electroconvulsive therapy relatively more. A persistently abnormal result after remission is a treatment-intensity signal, not a diagnosis.

Other probes surround the test. **The cortisol awakening response** is the rapid rise about 30 minutes after waking within a rhythm that then declines across the day; research uses its size, that rhythm's slope, stressor reactivity, and total output such as area under the diurnal curve or urinary cortisol. For localisation, **metyrapone** blocks cortisol synthesis and is normally followed by increased cortisol production, produces no such increase in adrenal or ectopic disease, and gives an augmented response in pituitary disease. Ovine CRH, corticorelin, and human CRH are used rarely, mainly to separate pituitary from ectopic ACTH sources during inferior petrosal sinus sampling, since pituitary ACTH production rises with CRH while ectopic does not; CRH after dexamethasone suppression also helps distinguish **pseudo-Cushing states**, as occur with alcoholism and some neuropsychiatric disorders, from true Cushing syndrome.

IN INDIA

Two divergences bite. The exposure history first: since prescribed steroid is the commonest cause of hyperadrenalism and is frequently taken here without a prescription in hand, one question about steroids under-detects it: ask separately about over-the-counter tablets, topical and combination preparations, injections given elsewhere, and unlabelled remedies, and count the inhaled route. Logistics next. The overnight test needs a timed late-evening dose and a timed morning sample, so without dependable outpatient phlebotomy it means a short admission or a failed test, and the shortcut taken instead, one convenient plasma cortisol, is what the source rules out. With phenytoin, barbiturates and carbamazepine in wide use and alcohol withdrawal a routine liaison context, nonsuppression here is more often the confounder than the disease.

The axis has also been a target, not only a marker. **The combined dexamethasone/CRF test** shows increased cortisol responses in major depression, but no protocol or threshold for it is available, so any claimed superiority over the plain test cannot be quantified. And interventions that suppress axis activity, including dexamethasone and the cortisol synthesis inhibitor ketoconazole, are sometimes used in more refractory depressive disorders, while attempts to develop antidepressants directly antagonising central CRH receptors have so far been unsuccessful. The pharmacology behind that failure belongs to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

Endogenous hypercortisolism depresses and impairs cognition while short exposure to prescribed steroid elevates, the direction set by duration rather than aetiology; onset clusters in the first fortnight; dose reduction is the treatment; and a suppression-test result means nothing until you name the question the test was asked.

4 · Adrenal insufficiency and the phaeochromocytoma differential

Two adrenal diseases, two psychiatric masks, one shared reason for being missed. Adrenal insufficiency reaches a psychiatrist as depression with fatigue in a patient who is quietly losing weight. Phaeochromocytoma reaches a psychiatrist as panic in a patient who is quietly hypertensive. Both are uncommon enough that nobody screens for them and consequential enough that missing either can kill.

Hold them as opposites: deficiency against excess, indolent against episodic, psychiatrically broad against psychiatrically narrow. In both the psychiatric label arrives first and the endocrine diagnosis late. Adrenal physiology itself belongs with the account of the hypothalamic-pituitary-adrenal axis elsewhere in this chapter. What is marked here is which presentation turns your attention to the adrenal gland, what you do next, and which of your own prescriptions can precipitate a catastrophe or corrupt the test you rely on.

40%	HIGH ACTH	TWO-THIRDS	24 h OR MORE
PHAEOCHROMOCYTOMAS THAT ARE PAROXYSMAL	PRIMARY, NOT SECONDARY, ADRENAL FAILURE	OF THE DIVIDED GLUCOCORTICOID DOSE, IN THE MORNING	FLUDROCORTISONE HALF-LIFE, SO NO DIVIDED DOSING

4.1 Where the lesion sits: primary against secondary failure

Everything downstream follows from the level of the lesion. **Primary adrenal insufficiency, or Addison disease**, comes from structural or functional lesions of the adrenal cortex itself; **secondary adrenal insufficiency** from lesions of the anterior pituitary or the hypothalamus.

The aetiological split is geographical and the most examinable fact here: autoimmune adrenal disease where resources are rich, **tuberculous adrenalitis** where they are not. The rest are short.

- Adrenalectomy, bilateral adrenal haemorrhage, and neoplastic infiltration of the adrenal glands.
- AIDS, which matters wherever tuberculosis matters.
- Inherited disorders of steroidogenic enzymes, and X-linked adrenoleukodystrophy.

Secondary disease behaves differently at the bedside because it generally presents more insidiously than the primary disorder, probably because mineralocorticoid biosynthesis is preserved: aldosterone is not under corticotropin control the way cortisol is, so the patient is spared salt-wasting. That is also why acute decompensation is more commonly associated with disorders of the adrenal than with disorders of the pituitary or hypothalamus.

AXIS	PRIMARY (ADDISON DISEASE)	SECONDARY
Site of lesion	Adrenal cortex, structural or functional.	Anterior pituitary or hypothalamus.
Commonest cause	Autoimmune adrenal disease where resources are rich; tuberculous adrenalitis where they are not.	Pituitary or hypothalamic dysfunction.
Mineralocorticoid	Lost, with the salt-wasting consequences.	Biosynthesis preserved.
Plasma corticotropin	High, from loss of normal glucocorticoid feedback inhibition.	Low.
Tempo	Insidious, but faster to declare itself.	Generally more insidious still.
Crisis liability	The commoner origin of acute decompensation.	Lower, for the same mineralocorticoid reason.

IN INDIA

Tuberculous adrenalitis is the most frequent cause of primary adrenal insufficiency in resource-limited regions. That changes the index of suspicion: the patient with treated or partially treated tuberculosis, or with HIV disease, who becomes apathetic, anorexic and postural is a different proposition from the same presentation in a European textbook.

4.2 The psychiatric phenotype of Addison disease

Onset is usually insidious, with progressive fatigue, weakness, anorexia, nausea, abdominal pain and weight loss, together with cutaneous and mucosal pigmentation, hypotension and

hypoglycaemia. The laboratory correlates are hyponatraemia, hyperkalaemia and metabolic acidosis. Read that list as a psychiatrist and every item is one you have written off before: fatigue, anorexia and weight loss as depression, abdominal pain as somatisation, postural symptoms as an antidepressant side effect.

How often the psychiatric features occur is genuinely unknown. Reports of neuropsychiatric symptoms in Addison disease are uncommon, although the true prevalence is uncertain because of probable under-reporting. No figure is available from the standard sources, and any percentage offered in an answer is invention. What the sources do describe is the shape.

STRATUM	FEATURES	WHAT IT MEANS FOR YOU
Predominant	Depression, reduced motivation and energy, and behavioural change.	What reaches psychiatry, mistaken for a primary mood disorder.
Commonest cognitive disturbance	**Memory dysfunction**.	Formal instruments belong to the Psychometry, Assessment and Descriptive Psychopathology notes.
Less common	Paranoid symptoms and delusions.	Not the modal picture; do not require psychosis before considering the diagnosis.
Rare but dramatic	Catatonia and self-mutilation.	The reason an unexplained catatonia gets a metabolic screen.
Heralding a crisis	Auditory and visual hallucinations, changes in conscious state, irritability, insomnia and nightmares.	A change in trajectory, not a new psychiatric diagnosis.
Established crisis	Frank delirium, coma and seizures.	An emergency; the delirium workup belongs to the Stroke, TBI, delirium and focal brain syndromes notes.

Hyponatraemia and metabolic acidosis may themselves contribute to the delirium and to the cognitive deficits, so part of what you are seeing is electrolyte disturbance rather than cortisol deficiency as such. Thresholds, mechanisms and the risks of correction are set out where this chapter deals with hyponatraemia and with the sodium-correction risk.

4.3 Adrenal crisis: recognising decompensation from the psychiatry ward

Acute adrenal insufficiency, the adrenal crisis, is the life-threatening end of the same disease. It is characterised by gastrointestinal symptoms of nausea, vomiting and abdominal pain, with dehydration, hyponatraemia, hyperkalaemia, weakness, lethargy and hypotension. Two things make it psychiatric. The prodrome is behavioural, in a patient who already carries a psychiatric label. And it can be iatrogenic, because it can follow abrupt withdrawal of glucocorticoids used at high doses or for prolonged periods: any patient on long-term steroids who stops them, or has them stopped during an admission, is exposed.

PITFALL

The insomnia, irritability and nightmares of an incipient crisis are read as psychiatric deterioration and answered with a sedative or a dose increase. What separates them is trajectory plus the somatic signs: falling blood pressure, vomiting, dehydration and clouding of consciousness are not features of a worsening mood or psychotic disorder. A patient with known adrenal insufficiency who becomes more disturbed needs observations and electrolytes before a psychotropic decision.

4.4 Localising the failure: corticotropin, stimulation and suppression

The single test that localises the lesion is the **plasma corticotropin immunoassay**. Its logic is feedback. Patients with primary adrenal insufficiency from intrinsic adrenal disease have high corticotropin levels because normal glucocorticoid feedback inhibition has been lost, whereas patients with secondary adrenal insufficiency from hypothalamic or pituitary disorders have low levels.

The same assay does the reciprocal job at the other end of the axis. In hypercortisolism, normal or high corticotropin levels are seen where the source is pituitary or ectopic, whereas low levels are invariably seen where the source is adrenal. The hypercortisolism syndromes are taught where this chapter deals with Cushing syndrome and steroid psychiatry. One caveat belongs to the assay itself: the specificity of these immunoassays for intact corticotropin can produce falsely low values in patients with ectopic secretion, so a low result in the wrong context is a reason to think again, not to close the case.

Function, as distinct from localisation, needs its own test. **Tetracosactrin testing**, in which an injection of a synthetic corticotropin analogue is given, assesses adrenocortical reserve. Dynamic endocrine tests divide cleanly: a stimulation test asks whether the gland can respond when driven, while **endocrine suppression tests** ask whether feedback is intact or secretion has become at least partly autonomous. Suppression of plasma cortisol by the synthetic glucocorticoid dexamethasone is incomplete in Cushing's syndrome, and suppression of serum growth hormone by an oral glucose load fails to occur in acromegaly. Two diseases, one principle.

The references name the stimulation test and its purpose but supply no usable dose, sampling interval or cortisol threshold, so none is given here. If asked for numbers, say what the test does and quote a threshold only from an endocrine source in front of you.

4.5 Management: replacement, crisis and the perioperative patient

LOCUS is binary. Chronic replacement is outpatient work shared with an endocrinologist, and the psychiatrist's contribution is adherence, recognising undertreatment and not creating a withdrawal crisis. Suspected acute adrenal insufficiency is emergency work: resuscitation with a steroid attached, where the correct psychiatric action is transfer, not titration on the ward. The perioperative patient sits between the two and is managed by protocol.

FOCUS shifts by phase. Acutely the target is circulatory collapse, dehydration and electrolyte derangement, with the neuropsychiatric features expected to follow the physiology. Short term,

stabilisation and step-down; longer term, a regimen that abolishes symptoms without imposing the harms of excess, plus the education that prevents the next crisis.

MODUS is almost entirely pharmacological, and the doses are the examinable content. For glucocorticoid replacement in chronic adrenal insufficiency, hydrocortisone was historically given at 20 to 30 mg per day, but most current guidelines now recommend the lower range of **15 to 20 mg per day**, based on estimates of normal endogenous cortisol production. The downward revision is the point: name both figures with the reason. The alternative agent for chronic replacement is cortisone acetate, inactive until it is converted to cortisol by 11-beta-hydroxysteroid dehydrogenase type 1, used at 25 to 37.5 mg per day. For mineralocorticoid replacement in primary adrenal insufficiency, **fludrocortisone acetate** is used at **0.05 to 0.2 mg per day**.

Both arms are titrated against clinical endpoints, not a target level. Mineralocorticoid dosing in primary adrenal insufficiency is based on resolution of electrolyte abnormalities and of postural hypotension, with reduction of salt craving as another valuable marker; the lower dose is used initially and titrated upward as required by blood pressure, plasma renin and the response to upright posture. Salt craving is the one a psychiatrist is most likely to elicit and least likely to ask about. Dosing frequency differs for a pharmacokinetic reason: fludrocortisone has a half-life of 24 h or more, so divided doses are not necessary, whereas the glucocorticoids are given in divided doses in an effort to mimic the normal diurnal rhythm of cortisol secretion, with two-thirds of the dose in the morning and one-third later in the day.

TRAP **Being asked when the second glucocorticoid dose is given.** The fractions are not in dispute: two-thirds in the morning, one-third later. The timing word is. The running text of the standard pharmacology reference places the smaller dose in the afternoon, while a table tip within the same reference states the evening. This is not pedantry: an evening glucocorticoid dose is precisely the one classically avoided for insomnia and for blunting the nocturnal cortisol nadir. Give the fractions, state the split as afternoon on the running-text reading, and name the variant rather than resolving it by assertion.

In the emergency, dose and route change entirely. **Critical illness-related corticosteroid insufficiency** reflects inadequate cortisol production, and can equally follow abrupt cessation of administered glucocorticoids. For that indication, high-dose intravenous hydrocortisone at 50 to 100 mg every 6 h, or a constant infusion at 10 mg per hour, is needed; prednisone at 1 mg per kg per day is an alternative. Bolus and infusion regimens are separate options and must never be blended into a hybrid figure. As the patient stabilises, the hydrocortisone dose may be decreased to 25 mg every 6 to 8 h, after which treatment proceeds as for chronic insufficiency. For perioperative cover in known adrenal insufficiency, hydrocortisone 100 mg parenterally every 8 h is given, and following surgery the dose is halved each day until routine maintenance levels are reached.

In suspected but unconfirmed acute adrenal insufficiency, **4 mg** of dexamethasone sodium phosphate can be substituted for hydrocortisone, because dexamethasone does not cross-react in the cortisol assay and will not interfere with measurement of cortisol either basally or in response

to the cosyntropin stimulation test. Treatment therefore starts before the diagnosis is secure without destroying the evidence, and failure to respond to cosyntropin in this setting is diagnostic of adrenal insufficiency. The same agent has a patient-held role: patients should carry dexamethasone 4 mg for intramuscular use in the event that severe nausea or vomiting precludes oral medication, and should then seek medical attention immediately.

GOOD TO REMEMBER

Sequencing matters when more than one axis has failed. In adrenal insufficiency associated with panhypopituitarism, glucocorticoids must be given before thyroid hormone is started, because thyroid hormone used in isolation accelerates metabolism of whatever endogenous cortisol remains and may precipitate an adrenal crisis. The temptation runs the other way, since the hypothyroidism is often noticed first. Thyroid replacement is taught where this chapter deals with hypothyroidism and hyperthyroidism; what belongs here is the order.

4.6 The psychiatric symptoms respond to the steroid, not to a psychotropic

The psychiatric literature is unusually strong on this point. The neuropsychiatric symptoms of Addison disease, including those of adrenal crisis, appear largely to resolve once treatment with adequate doses of corticosteroids is commenced, and the use of other psychotropic agents is rarely indicated.

RULE **In diagnosed adrenal insufficiency, persistent psychiatric symptoms are a replacement-dose question before they are a psychotropic question.** Recent data indicate that higher corticosteroid doses reduce the occurrence and severity of neuropsychiatric symptoms, particularly depressive symptoms. The first move when mood and motivation fail to recover is to ask whether replacement is adequate, with whoever manages the endocrine disease, rather than to layer a psychotropic on top.

Hold that direction alongside its mirror image. In replacement, more steroid means fewer psychiatric symptoms; in supraphysiological exposure, which this chapter covers under steroid-induced psychiatric disturbance in Cushing syndrome and in therapeutic corticosteroid use, more steroid means more psychiatric risk. They are not in conflict: they describe different halves of one dose-response curve, correcting a deficiency and imposing an excess. No milligram figure attaches to "higher doses" in the source, so it is a direction of adjustment and not a licence to name a number.

4.7 Phaeochromocytoma: the panic phenocopy

Phaeochromocytomas are catecholamine-secreting tumours that most commonly originate in the chromaffin cells of the adrenal medulla, though they may rarely arise from similar cells in sympathetic ganglia. Familial forms are associated with **multiple endocrine neoplasia syndrome type IIa and type IIb**, and with von Recklinghausen neurofibromatosis. Those are the two associations the standard psychiatric reference names; other syndromic and mutational associations circulate in endocrine practice but are not asserted here.

Catecholamine secretion may be continuous or sporadic, and when sporadic it produces characteristic paroxysmal symptoms: headache, profuse sweating, palpitations, Raynaud

phenomenon, tremor, nausea, vomiting, and abdominal and chest pain. The clinical signs are hypertension, pallor, postural hypotension, and the signs of chronic hypertension such as retinopathy. Read that against the criteria for a panic attack and the overlap is near complete: this tumour is a psychiatric problem before it is a surgical one.

MNEMONIC

Palpitations, Headache, Sweating

The triad of palpitations, headache and profuse sweating is the most sensitive and specific presentation for phaeochromocytoma. All three are symptoms a panicking patient may report, so the triad narrows the field without settling it. The source asserts its discriminating power without attaching sensitivity or specificity values, so quote it qualitatively.

40% of phaeochromocytomas present with paroxysmal symptoms, and the paroxysm may present as a phenocopy of anxiety symptoms, particularly a panic episode, and is often diagnosed as such. The precipitants are informative because several are shared with panic and one is not: episodes may occur spontaneously or be precipitated by exercise, postural change, raised intra-abdominal pressure, or emotional excitement or shock. Raised intra-abdominal pressure has no psychiatric counterpart and is worth asking about. Severity varies between episodes, and anxiety may persist for some time after the attack, so residual anxiety after a paroxysm does not argue for a primary anxiety disorder.

The negative finding narrows the differential. Psychosis and cognitive deficits have not been described in phaeochromocytoma, although there are several case reports of a secretory tumour causing relapse of a previously treated or quiescent psychotic disorder. This is an anxiety-spectrum problem: a first-episode psychosis or a progressive cognitive syndrome is not a phaeochromocytoma presentation.

4.8 Confirming it, treating it, and the psychotropics that intervene

Diagnosis is made by detecting increased urinary secretion of catecholamines or catecholamine metabolites. The psychiatric reference goes no further: no assay named, no sensitivity or specificity given, no cut-off set, so none is stated. Locating the tumour and excising it surgically, under alpha-adrenergic blockade, is effective in resolving most anxiety symptoms if the paroxysmal episodes are terminated. No alpha-blocking agent is named and no preoperative interval is given in these sources, so the teachable content is the sequence and not a drug or a schedule.

TRAP **Reaching for a beta blocker first in the sweating, tremulous, palpitating patient.** Propranolol is a familiar prescription for tremor and performance anxiety, and the reflex is strong. The monograph warning is categorical: do not use propranolol in phaeochromocytoma unless alpha blockers are already being used. The examinable answer is the order: alpha before beta. The source states the prohibition without explaining it, so the safe answer names the sequence and stops there rather than embellishing with a mechanism.

Two of your own prescriptions can corrupt or contraindicate this workup. False-positive results may occur in screening tests for phaeochromocytoma in patients being treated with prazosin, and if an elevated urinary vanillylmandelic acid is found, prazosin should be discontinued and the patient retested after a month. Prazosin is prescribed for nightmares in post-traumatic stress disorder, so this collides with real practice: the traumatised patient with hyperarousal and autonomic symptoms is precisely the one in whom someone orders a catecholamine screen. The second group is contraindicated outright, and the list is a psychiatric one.

- Atomoxetine: do not use if the patient has phaeochromocytoma or a history of phaeochromocytoma, or a severe cardiovascular disorder that might deteriorate with clinically important increases in heart rate and blood pressure.
- Isocarboxazid, phenelzine and selegiline: do not use if the patient has phaeochromocytoma.
- Amisulpride, sulpiride, flupenthixol and pipothiazine: do not use if the patient has phaeochromocytoma.

Atomoxetine is the clinically live entry, a mainstay of non-stimulant attention deficit hyperactivity disorder prescribing in Indian practice. The monoamine oxidase inhibitor entries are mechanistically unsurprising. The last group is the one candidates forget, because the substituted benzamides and thioxanthenes are not drugs anyone associates with catecholamine physiology. Choice among antipsychotics otherwise belongs to the Schizophrenia: management, antipsychotics and adverse effects notes; what belongs here is that this diagnosis removes four of them from the table.

> Adrenal insufficiency presents as depression, apathy and memory failure that respond to adequate corticosteroid replacement rather than to a psychotropic; phaeochromocytoma presents as panic, and in a patient who might have one, alpha blockade always precedes beta blockade.

5 · Hypothyroidism and hyperthyroidism in psychiatry

No gland is asked about more often than this one. The examinable content is not that thyroid disease produces psychiatric symptoms but the shape of the association: which symptoms arrive before the physical features, and what is still there a year after the gland has been corrected. Almost every error here is a caricature. The underactive gland is supposed to produce a slowed, depressed, forgetful patient and the overactive gland a bright, fast, manic one. The first caricature is roughly right; the second is wrong, and the prevalence figures say so plainly enough that examiners build questions on it.

The two diseases are not mirror images. What genuinely inverts between them is the timing relative to the physical illness and the degree to which the biochemistry predicts the psychiatry.

UP TO 50%	APPROXIMATELY 40%	UP TO 30%	2 TO 5%
COGNITIVE DEFICITS IN HYPOTHYROIDISM	DEPRESSION IN HYPOTHYROIDISM	DEPRESSION IN HYPERTHYROIDISM	MANIC SYMPTOMS IN HYPERTHYROIDISM

5.1 Two causal lists, and the test logic that anchors both

Hypothyroidism in adults is overwhelmingly a disease of the gland itself. The commonest cause is **primary autoimmune hypothyroidism related to antithyroid antibodies**, with the remainder made up largely of treatment given for hyperthyroidism, drug-related effects, and iodine deficiency. Causes above the gland, meaning hypothalamic or pituitary dysfunction, account for less than 5% of cases. That proportion carries the interpretive logic of the thyroid function test. Because the great majority is primary, loss of negative feedback drives thyroid-stimulating hormone up while the peripheral hormone falls, so in overt hypothyroidism a raised thyroid-stimulating hormone with a low free thyroxine is the expected pattern. The rare central case breaks it: a low free thyroxine with a thyroid-stimulating hormone that is low or merely unremarkable is a pituitary problem until proven otherwise, and needs the rest of the anterior pituitary examined rather than a prescription for replacement.

Hyperthyroidism has a broader causal list because several different mechanisms produce the same excess. The excessive production of thyroid hormone may arise from a toxic multinodular goitre, from a single functioning adenoma, or from a thyroid stimulator such as the **thyroid-stimulating antibody in Grave disease**. Exogenous thyroid hormone reproduces the picture, as do disorders of thyroid hormone storage consequent to autoimmune thyroiditis. Two matter disproportionately in a psychiatric clinic: the patient taking thyroid hormone for weight control or as an antidepressant adjunct, and the patient whose autoimmune thyroiditis is releasing stored hormone before it burns out into hypothyroidism, so that one gland delivers both syndromes in sequence.

Two boundaries belong to neighbouring material. The numeric definitions that separate subclinical from overt hypothyroidism sit with this chapter's account of subclinical thyroid disease, autoimmunity and Hashimoto encephalopathy, and they are not symmetrical: the hypothyroid side carries an agreed thyroid-stimulating hormone band, the hyperthyroid side has no equivalent, and a candidate who constructs one by mirror image is inventing it. Lithium-induced thyroid dysfunction is taught with the lithium and thyroid material; lithium as a drug, its pharmacokinetics and its non-thyroid organ effects belong to the Mood disorders: pharmacotherapy notes.

5.2 Myxoedema: the examination that makes the diagnosis

The symptoms include fatigue, lethargy, weight gain, constipation, cold intolerance, stiffness and cramping of muscles, hair loss, cognitive slowing and depression. Read that list as a psychiatrist and it is indistinguishable from a moderate depressive episode with somatic features, which is exactly why the examination decides the question.

- General and cardiovascular: hypothermia and bradycardia.

- Skin, hair and face: dry skin, sparse hair and periorbital swelling.
- Mouth and voice: thickening of the tongue, with coarsening and deepening of the voice.

The sign that settles it at the bedside is the **characteristic prolonged relaxation phase of the deep tendon reflexes**, which needs no laboratory and is the one finding a depressed patient cannot produce. Taken together this constellation is what is meant by **myxoedema**.

GOOD TO REMEMBER

The free thyroxine value has no number attached to it in these notes on purpose: report it against the reporting laboratory's own reference interval, since assays and units differ and a remembered range is a good way to call a normal result low. The direction of change, read together with the thyroid-stimulating hormone, carries the diagnosis.

5.3 The hypothyroid ladder: cognition, then mood, then anxiety

Cognitive deficits are the most common neuropsychiatric feature, occurring in up to 50% of cases, with impairment of psychomotor speed, memory and visual-perceptual skills. Difficulty with constructional skills and reduced performance on trail-making and maze tasks point to prominent executive deficits as well, so this is not a pure amnestic picture but a diffuse slowing with an executive edge. How cognition is examined belongs to the Psychotherapies, clinical assessment, MSE and nosology notes, and the dementia syndrome and its workup to the Dementias and neurocognitive disorders notes.

Depression is only slightly less common than cognitive disturbance, reported in approximately 40% of patients, and presents most often as low mood, fatigue, anhedonia, reduced concentration and hypersomnolence. Generalised anxiety symptoms are described in up to 30% of patients, strongly correlated with the depressive symptoms but not with biochemical severity. The mood disturbance is not always unipolar: it may take a **rapid-cycling form**, which is the reason thyroid function is checked in a patient whose bipolar illness has accelerated rather than only in one who looks flat.

MNEMONIC

Think, then feel, then fear, then break

Hypothyroidism descends that ladder in order of frequency: cognitive impairment is commonest, depression next, generalised anxiety next, and frank psychosis is rare. The same four presentations occur in thyrotoxicosis in a different order, so it is the ordering rather than the list that is worth memorising.

The severity of the cognitive disorder is correlated with the degree of biochemical abnormality, and although largely corrected by a return to a euthyroid state, some deficits may remain, particularly in older adults or in those with reduced cognitive reserve. Depression behaves differently: it is less closely related to the severity of the biochemical hypothyroidism. Severe disturbance of consciousness, including coma and delirium, may be encountered, but may be associated with other metabolic changes accompanying the hypothyroid state, such as

hyponatraemia: that sodium pathway is taken up where this chapter deals with SIADH and hyponatraemia, and delirium as a syndrome belongs to the Stroke, TBI, delirium and focal brain syndromes notes.

RULE **Biochemical severity predicts the cognitive picture, not the mood picture, in either direction of thyroid dysfunction.** Cognitive severity tracks the degree of biochemical abnormality in hypothyroidism; depressive severity does not, in either hypothyroidism or hyperthyroidism. A near-normal set of results therefore never excludes a thyroid-driven depression, and a grossly abnormal set never predicts how depressed the patient will be. Anxiety splits between the two: it follows the depression in hypothyroidism and more often follows the thyrotoxic severity in hyperthyroidism.

5.4 Myxoedema madness: the eponym that outlived its epidemiology

Psychotic symptoms in hypothyroidism, including paranoid ideas, misidentification, visual and auditory hallucinations and thought disorder, were originally thought to be common and were described as **myxoedematous madness**, but they likely occur in **less than 5%** of all patients with hypothyroidism and tend to emerge after the onset of the physical symptoms. So a first-episode psychosis in a physically well patient is not myxoedema madness by default.

Indian consultation-liaison practice adds a presentation the standard psychiatric account does not name. Subclinical to moderate hypothyroidism is often associated with depressive symptoms, while severe hypothyroidism can give rise to psychosis, including myxoedema madness, delirium and **catatonia**. That completes a severity gradient: mild and moderate disease is a mood problem, severe disease is a brain-failure problem, and catatonia sits at the severe end alongside the psychosis and the delirium rather than being a separate curiosity.

TRAP **Treating myxoedema madness as a common and early presentation of hypothyroidism.** The historical literature described it as frequent and candidates reproduce that framing. The corrected version is that psychosis occurs in under one in twenty patients with hypothyroidism and follows rather than precedes the physical features. Write the eponym, then correct it in the same breath; that sequence is what earns the mark.

IN INDIA

Two divergences are real here. First, the causal list is weighted differently: autoimmune disease still leads, but iodine deficiency remains a live entry in a way it is not in most Western practice, so a hypothyroid patient from an iodine-deficient area is not automatically an autoimmune one and the dietary and geographic history is worth taking. Second, Indian consultation-liaison teaching explicitly places catatonia within the severe-hypothyroid picture: where a catatonic patient may reach psychiatry before any medical workup, thyroid function belongs in the first set of investigations, and calling a catatonia idiopathic before it has been checked is a decision, not an omission.

5.5 Why the hypothyroid brain becomes depressed

The origin of depression in hypothyroidism appears to relate to the role of thyroxine in serotonergic transmission: reduced thyroid input reduces serotonergic tone and lowers the

threshold toward the development of depressive symptoms, while thyroid replacement restores central serotonin activity in concert with improvement in the depressive symptoms. The same mechanism is thought to underpin the adjunctive antidepressant effect of thyroxine, which is the bridge to thyroid hormone augmentation in resistant depression. That strategy, with its agents and doses, is taught where this chapter deals with lithium, thyroid dysfunction and thyroid hormone augmentation, and none of it should be reconstructed from the physiology given here.

The imaging tells the same story at the level of the whole brain. Functional neuroimaging has demonstrated global hypometabolism, with more specific areas of hypoperfusion in the posterior cingulate, the insula, the fusiform gyrus, and the right parieto-occipital and primary motor cortex. The global finding fits the clinical picture better than any regional one: a diffusely slowed brain produces diffusely slowed cognition, which improves as the metabolic state is restored. The wider principles of functional imaging and neurotransmitter measurement belong to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

5.6 Hyperthyroidism: depression first, mania last

Here the caricature fails completely. Depressive symptoms are the most common psychiatric feature of hyperthyroidism, occurring in up to 30% of patients, and frequently occur prior to the onset of the other physical features. This includes lowering of mood together with neurovegetative disturbance such as insomnia, reduced libido, weight loss and fatigue. The discriminator is one symptom. As distinct from depression, appetite is invariably increased in hyperthyroidism, so **increased appetite with weight loss**, in a patient who also reports insomnia and fatigue, points at the thyroid until the test says otherwise.

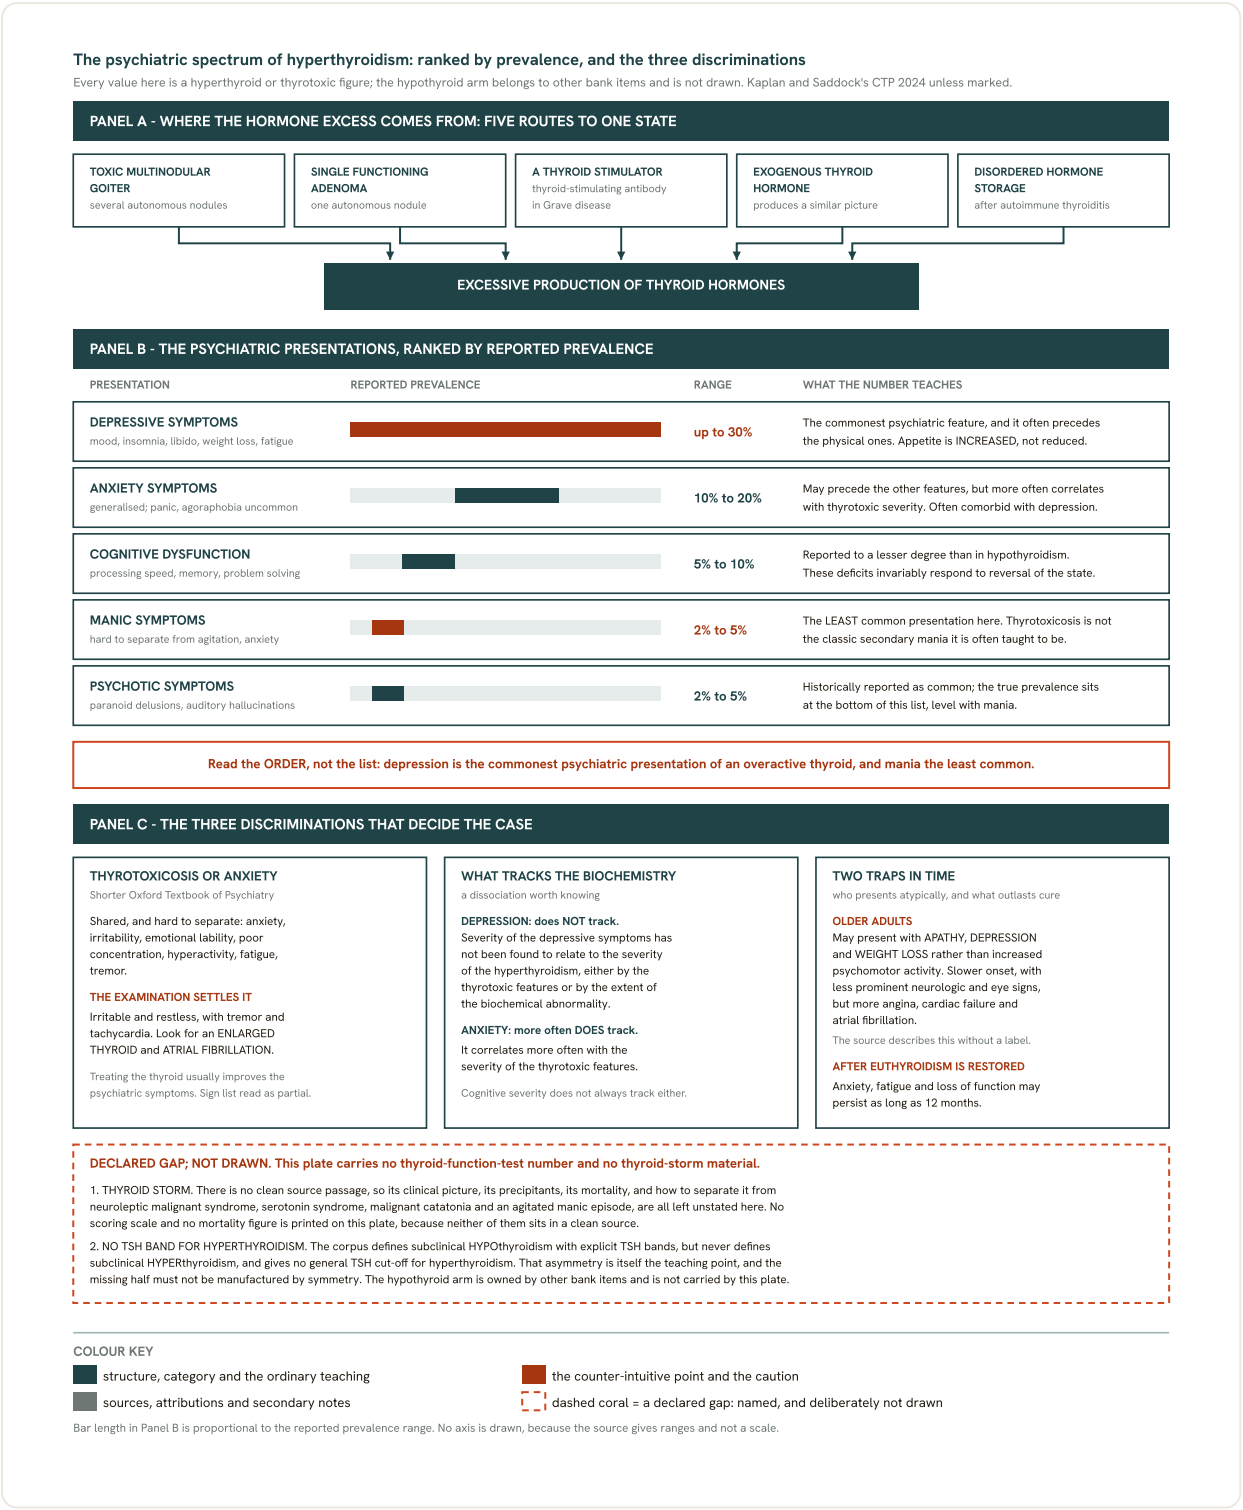


Fig 33.3 **The psychiatric spectrum of hyperthyroidism, ranked by reported prevalence.** The ordering is the teaching: depressive symptoms are the commonest psychiatric feature of an overactive thyroid and frequently precede the physical ones, while manic and psychotic symptoms sit at the bottom of the list, which contradicts the folk teaching that thyrotoxicosis is a classic secondary mania. Appetite is invariably increased, and that is the discriminator from a primary depressive episode. Depressive severity does not track the biochemistry; anxiety more often does. In older adults the presentation may be apathy, depression and weight loss rather than overactivity, an entity the source describes without naming. The hypothyroid arm of the spectrum and every thyroid-function-test band are declared gaps for this item and are not drawn. (Kaplan and Saddock's CTP 2024; Shorter Oxford Textbook of Psychiatry for the anxiety-disorder differential and the examination signs, where the sign list is read as partial.)

The remaining presentations fall in a strict order of frequency below depression.

- Anxiety presenting as generalised anxiety is common, with a prevalence of between 10% and 20%, but panic and agoraphobia are relatively uncommon.
- Cognitive dysfunction is reported in **between 5% and 10%** of patients with thyrotoxicosis, to a lesser degree than in hypothyroidism.
- Manic symptoms are less common, with a prevalence of **between 2% and 5%**, and may be difficult to distinguish from psychomotor agitation and anxiety; psychotic symptoms, including paranoid delusions and auditory hallucinations, were historically reported as common but have a true prevalence of between 2% and 5%.

Timing separates anxiety and depression from everything else. Anxiety, like the depressive symptoms, may occur prior to the other features of hyperthyroidism, but more often correlates with the severity of the thyrotoxic features, and anxiety and depressive features are frequently comorbid. Depression does not behave that way: its severity has not been found to relate to the severity of the hyperthyroidism as measured by the subsequent thyrotoxic features or the extent of the biochemical abnormality.

TRAP **Naming thyrotoxicosis as a classical cause of secondary mania.** It is taught that way in wards and it is the wrong end of the list: mania is the least common of the four psychiatric presentations of hyperthyroidism, sitting alongside psychosis at the bottom, while depression sits at the top at up to three in ten patients. The useful version is the reverse: suspect thyrotoxicosis in the depressed patient with weight loss and a rising appetite, not principally in the elated one.

5.7 The two clinical mistakes: the anxious young patient and the apathetic old one

The first mistake is diagnostic. Hyperthyroidism may present with anxiety, irritability, emotional lability and difficulty in concentrating, and these, together with hyperactivity, fatigue and tremor, can make the differential diagnosis from an anxiety disorder difficult. The resolution is physical rather than psychopathological. The patient may be irritable and restless with tremor and tachycardia, and examination may reveal characteristic signs such as an enlarged thyroid and atrial fibrillation; that list of signs is reliable but not necessarily exhaustive, so a normal thyroid on palpation does not close the question and the biochemistry still decides it. Treatment of the thyroid dysfunction usually results in improvement of the psychiatric symptoms, which is the retrospective proof but no help at the first consultation.

The second mistake is a failure of expectation in older patients. Older adults with thyrotoxicosis may present with predominantly apathy, depression and weight loss rather than increased psychomotor activity, typically with a slower onset and with less prominent neurologic and ophthalmologic features, while cardiovascular events such as exacerbation of angina, cardiac failure and atrial fibrillation are more prevalent. This is the presentation conventionally labelled apathetic thyrotoxicosis, and every element of it points away from the diagnosis. In that patient the thyroid is disclosed by the heart rather than by the eyes.

AXIS	HYPOTHYROIDISM	HYPERTHYROIDISM
Depression	Approximately 40%; low mood, fatigue, anhedonia, reduced concentration, hypersomnolence.	Up to 30%; the commonest psychiatric feature, with appetite increased rather than reduced.
Anxiety	Up to 30%; generalised.	Between 10% and 20%; panic and agoraphobia relatively uncommon.
Cognition	Up to 50%; psychomotor speed, memory, visual-perceptual and executive.	Between 5% and 10%; slow processing, immediate memory, higher-level problem solving, or frank delirium.
Psychosis	Less than 5%.	Between 2% and 5%.
Elevated mood	Mood disorder may take a rapid-cycling form.	Manic symptoms between 2% and 5%; the least common presentation.
Timing versus physical features	Psychosis tends to emerge after the physical symptoms.	Depression frequently precedes the physical features, and anxiety may also precede them.
Tracks biochemical severity	Cognitive severity yes; depression no; anxiety no.	Depression no; anxiety more often yes; mild attention and concentration deficits not always.
After euthyroid state restored	Mood symptoms may reverse but neuropsychiatric problems may persist; some cognitive deficits may remain.	Cognitive deficits invariably respond; anxiety, fatigue and loss of function may persist as long as 12 months.

5.8 Management: correcting the gland, and knowing what correction will not fix

LOCUS. Most of this is outpatient work shared with a physician, where the psychiatrist contributes recognition and sequencing rather than prescription. Inpatient care is for the severe end: the psychosis, delirium and catatonia of severe hypothyroidism, and the patient whose disturbance of consciousness makes safe outpatient correction impossible. Where thyrotoxicosis becomes an acute medical crisis the locus is a medical unit, and the psychiatrist's job is to keep the differential of an agitated, unstable patient open rather than closing it onto a psychiatric label. The authorities here are Kaplan and Sadock's Comprehensive Textbook of Psychiatry for the neuropsychiatric epidemiology and its treatment, the Shorter Oxford Textbook of Psychiatry for the differential against primary anxiety and mood disorder, and Indian consultation-liaison guidance for the severity gradient and the catatonia entry.

FOCUS changes with phase. Acutely the targets are safety, the psychotic or catatonic presentation, and the medical instability that produced them. Over the short term the target is the euthyroid state itself, and this is where restraint earns most: the depressive syndrome seen in hyperthyroidism rarely requires treatment other than that used to restore the euthyroid state. Over the mid term the target is what fails to resolve, since anxiety, fatigue and loss of function may persist as long as **12 months** after a euthyroid state has been achieved. Long term the target

is residual cognition, because replacement therapy may reverse the mood symptoms while neuropsychiatric problems persist, with older adults and those with reduced cognitive reserve the group in whom that residue is most likely.

MODUS begins with the endocrine treatment, which is the effective psychotropic in both diseases, and adds psychiatric treatment only where a specific indication survives it. The psychotic symptoms of hypothyroidism also respond to appropriate thyroid hormone treatment within days to weeks, but rapid titration of hormone doses may exacerbate the psychosis, and careful addition of a low dose of antipsychotic to thyroxine has been reported to be well tolerated and to result in earlier remission of the psychosis. Two instructions sit inside that: go slowly with the hormone in a psychotic patient, and use the antipsychotic alongside it rather than instead of it. No numeric dose of either agent is stated in the source and none should be supplied from memory; the choice of antipsychotic agent and its receptor pharmacology belong to the Schizophrenia: management, antipsychotics and adverse effects notes. The psychological and social modes are load-bearing: explaining to a patient depressed for months that the treatment is for a gland, and warning patient and family that anxiety and fatigue may outlast the biochemistry by many months, prevents the demoralisation that follows a correction which does not feel like a cure.

PITFALL

The residual state after successful treatment of hyperthyroidism gets misdiagnosed twice over. Anxiety and fatigue persisting months after the thyroid is corrected are read as a new primary anxiety disorder and treated as one, or read as inadequate control and used to justify further meddling with the thyroid. These features can persist for up to a year in a properly euthyroid patient, so the correct action is usually to name the timeline, keep reviewing, and not start a new illness on the strength of it.

Depression is the commonest psychiatric presentation of BOTH thyroid diseases, mania is the least common presentation of hyperthyroidism, and biochemical severity tracks the cognitive deficit but never the mood.

6 · Lithium, thyroid dysfunction and thyroid hormone augmentation

Two questions live here and they run in opposite directions. The first is what lithium does to a healthy thyroid, and how to hold a patient on the drug carrying their mood disorder while that thyroid quietly fails. The second is the mirror: what thyroid hormone does to a mood disorder in a thyroid that works. Lithium as a drug belongs to the Mood disorders: pharmacotherapy notes; interpreting thyroid function tests is set out where this chapter deals with hypothyroidism and hyperthyroidism in psychiatry. What is added here is lithium and thyroid function taken together: what those results mean when lithium is why you ordered them.

1.5% PER YEAR	8% TO 19%	UP TO 20%	3 TO 4-FOLD
NEW HYPOTHYROIDISM, ANNUAL INCIDENCE	PREVALENCE, OVERT DISEASE	RISK IN MIDDLE-AGED WOMEN	THYROID LITHIUM VERSUS PLASMA

6.1 Why the thyroid, and by what mechanism

Lithium concentrates in the organ it disables, by the route iodide takes. The **Na+/I- symporter** carries iodide into the follicular cell, where thyroid peroxidase oxidises it and then iodinates selected tyrosine residues on thyroglobulin. Lithium is a substrate for that transporter, so it accumulates: levels in the gland run three to four times higher than in plasma, which is why organ-specific effects appear at ordinary plasma levels. Its actions fall into three groups.

- Release is inhibited: T4 and T3 are held rather than secreted.
- Thyroglobulin conformation is altered, so coupling of its iodinated tyrosine residues to form T4 and T3 is inhibited.
- Binding of T4 and T3 to hypothalamic receptors is altered, downregulating expression of genes for certain receptor isoforms.

MNEMONIC

Release, Reshape, Receptor

Two of the three sit inside the gland, one above it, which is why the picture is not always a clean primary hypothyroidism.

Reassuringly, these effects are largely reversible and lithium is not associated with the development of antithyroid antibodies, though an occasional patient has persistent hypothyroidism after lithium is stopped, for reasons lying elsewhere.

TRAP **Writing that lithium blocks iodine uptake into the thyroid.** Widely reproduced in Indian examination material, and with a real source: Indian consultation-liaison guidance states that lithium interferes with iodine uptake, iodination of tyrosine, thyroglobulin structure and hormone release. The Lithium Handbook contradicts the first element flatly: lithium alters neither iodide uptake nor efflux, being a substrate rather than a blocker. They agree on everything downstream. Name the agreed mechanisms, then flag uptake as disputed rather than asserting either version as settled.

6.2 How common, in whom, and the numbers that get blended

At least four figures for lithium-induced thyroid disease circulate, all of them for hypothyroidism, and each is correct in its own frame. The tightest concerns new disease over time: in long-term follow-up the annual incidence of new hypothyroidism is **approximately 1.5%**, with females affected at three times the rate of males. That is an incidence, new cases per year of exposure: the question a patient asks at initiation.

A prevalence answers something else. The prevalence of **overt hypothyroidism** in lithium-treated patients runs from 8% to 19%, a band restricted to overt disease and so excluding the larger subclinical group. A third is subgroup-bound: the Maudsley Prescribing Guidelines note that in middle-aged women the longer-term risk may reach 20%. Indian consultation-liaison teaching carries a fourth and much broader band, overlapping all of these without matching any.

The recognised predictors group usefully.

- Fixed patient factors: female sex, a predisposition to autoimmune hypothyroidism, family history.
- Age older than 50 years.
- Treatment-related: duration of lithium treatment, and weight gain of more than 5 kg.

The last two are observable without any extra test: a patient who has gained substantial weight over two years of lithium is one whose next TSH deserves attention. For proportion, hypothyroidism does not rank among the leading somatic reasons patients discontinue lithium.

TRAP **Reciting "20%" as the risk of hypothyroidism on lithium.** A candidate who cannot say whether the number is annual or cumulative, overt or subclinical, or in whom, will be taken apart in a viva. Do not average the four. Name the quantity, the population, the period.

6.3 Reading the TSH: reference range and the argument inside it

The normal TSH range is **0.35 to 4.5 mIU/l**, and **subclinical hypothyroidism** occupies 4.5 to 9.9 mIU/l. Resist hardening those edges: the same source records genuine debate about the upper limit, hedges it elsewhere as typically 4.5 or 5.0, and notes the argument that for younger females the upper limit of normal might be as low as 2.5 mIU/l. Subclinical thyroid disease as an entity is developed where this chapter deals with subclinical thyroid disease and Hashimoto encephalopathy.

The source argues for a flexible rather than threshold-driven approach: a patient with a TSH of 4.0 mIU/l and significant hypothyroid symptoms may do better with a TSH nearer 1.00 mIU/l.

Cohort evidence agrees. The investigators who linked lower mean free T4 to more affective episodes and greater depression severity then found that lithium-treated bipolar patients who required treatment for depression had a significantly higher adjusted mean TSH, 4.4 mIU/l, than those who did not, 2.4 mIU/l. Both are group means, not thresholds, and both sit inside the normal range: an argument for a lower threshold to intervene, not for redefining normality.

GOOD TO REMEMBER

Fatigue, weight gain, cognitive slowing and low mood in a lithium-treated patient are the presenting complaints of both a depressive relapse and an underactive thyroid. The discrimination is laboratory, not bedside. Order the TSH before you rework the mood-stabiliser regimen, not after it fails.

6.4 Baseline work-up, and how often to repeat it

Before lithium is started, renal, thyroid and cardiac function should be checked, with estimated glomerular filtration rate, urea and electrolytes and thyroid function tests as the minimum, and a baseline calcium is also helpful. Only a pre-treatment calcium will later separate drug effect from pre-existing disease; the calcium-related syndromes are developed where this chapter deals with hyperparathyroidism and hypoparathyroidism. Once established, two independent sources agree on the interval at which lithium and thyroid function should be re-checked.

SOURCE	ESTABLISHED THERAPY	WHAT IT ADDS
The Lithium Handbook	Serum calcium, TSH and blood pressure every 6 months, with lithium level and estimated glomerular filtration rate on the same schedule.	Frequency changes if an abnormality is found or treatment started.
Maudsley Prescribing Guidelines, fifteenth edition	Plasma lithium, estimated glomerular filtration rate, urea and electrolytes and thyroid function tests every 6 months.	More frequent testing with interacting drugs, in older patients, or in established chronic kidney disease.
Geriatric psychiatry text (dissenting)	Thyroid parameters every 6 to 12 months once established on a stable dose.	Check before initiation and again after every dose adjustment.

Two clean sources against one is why the figure to teach is **every 6 months**; the dissenter is a geriatric source whose surrounding passage is textually damaged. Show the disagreement anyway: older-adult practice may legitimately differ, and the dissenter adds one element the others omit, a re-check after every dose adjustment, contradicted by nothing.

No source here says when the TSH should first be repeated after initiation: the handbook advises discussing when to repeat it without naming an interval, the Maudsley argues for more tests in the first year without a figure. The established-therapy cadence is well grounded; the initiation cadence is local protocol.

6.5 How far to investigate: antibodies, imaging and a live disagreement

Two respected sources take opposite positions on autoantibody testing, both defensible because their scopes differ. The Lithium Handbook rests on an empirical finding. long-term lithium treatment is not associated with an increased prevalence of autoantibodies to thyroid peroxidase, thyroglobulin or the TSH receptor compared with age-matched and gender-matched peers, and carries no increased risk of thyroid tumours. If lithium does not generate autoimmunity, measuring autoimmunity does not measure lithium. Hence there is no indication for routine thyroid ultrasound or autoantibody testing, and any laboratory measure beyond TSH should be dictated by the clinician prescribing L-thyroxine and by the clinical course.

The Maudsley position is narrower. Discussing middle-aged women, it records that a case has been made for testing thyroid autoantibodies in this group before starting lithium, to estimate risk better, and for performing thyroid function tests more frequently during the first year. That is risk stratification in a defined subgroup, not routine surveillance, so the two do not collide head-on. But the reconciliation is only partial: the handbook's evidence undercuts the risk-stratification rationale as much as the surveillance one. The risk-factor list carries a related tension, naming a predisposition to autoimmune hypothyroidism while the handbook finds no lithium-induced rise in antibody rates; both hold if kept apart, since a patient may bring autoimmune vulnerability to lithium without lithium creating it.

PITFALL

Reading a negative thyroid autoantibody panel as reassurance that this patient will not become hypothyroid on lithium. Evidence that lithium does not raise antibody titres is evidence about mechanism, not a negative predictive value for the individual. Antibody status does not replace the TSH, nor licence a longer monitoring interval.

6.6 Managing hypothyroidism without stopping the lithium

LOCUS. Outpatient almost without exception: the test is a TSH, the treatment a tablet, so neither admission nor endocrinology referral is needed. Shared care matters at two points only, when the TSH is low rather than high, and in pregnancy.

FOCUS. Before treatment, a baseline that lets later change be attributed correctly. Acutely, rarely anything. Short and medium term, the symptom overlap. Long term, continuity of the mood-stabiliser treatment, where nearly all the harm in lithium-induced thyroid disease is done.

MODUS. Three sources govern: the Maudsley Prescribing Guidelines for monitoring, The Lithium Handbook for the treatment decision, and the American Thyroid Association pregnancy guidelines for the obstetric adjustment. **L-thyroxine** is the treatment of choice rather than lithium discontinuation, started on the TSH. A prior history of hypothyroidism is no reason to deprive a patient of lithium; where it precedes lithium, treatment begins as lithium commences.

RULE **Hypothyroidism never justifies stopping lithium, new-onset included.** **Professor Michael Gitlin**, in a 2016 review of lithium's adverse effects quoted approvingly in The Lithium Handbook, calls this the most important clinical rule in the area. You would be trading a fully treatable endocrine deficiency for the best-evidenced prophylaxis in bipolar disorder, and the relapse risk that follows discontinuation. Supplement the thyroid, keep the lithium.

The sources that insist L-thyroxine is the treatment of choice give no starting dose, no titration increment, no titration interval, no TSH target, and no weight-based figure appears anywhere in this material. Prescribe from local thyroid guidance and the clinician managing the replacement; do not reconstruct a dose from memory.

Pregnancy is the one place a specific figure is grounded. In a woman already on thyroid supplementation, high-quality evidence supports increasing the levothyroxine dose by 20% to 30% on a positive pregnancy test, with follow-up by her treating clinician: a proportional adjustment to a dose already taken, not a starting dose. Starting thresholds in pregnancy are separate: a strong recommendation, on low-strength evidence, to treat subclinical hypothyroidism above 10 mIU/l; a weak one between the pregnancy-specific upper limit and 10 mIU/l; and in women whose thyroid peroxidase antibodies are diagnostic of Hashimoto's thyroiditis, treatment above that pregnancy-specific limit, with weak consideration above 2.5 mIU/l while still within it. That last figure is pregnancy and antibody-specific, unrelated to the debate about the upper limit of normal in younger women, which uses the same number for a different population.

IN INDIA

Two divergences. First, the recommendation is contested here. Indian consultation-liaison guidance offers switching to a different mood stabiliser where clinically significant hypothyroidism occurs or subclinical hypothyroidism persists for a defined period, and names T3 rather than L-thyroxine as the thyroid supplement. The Lithium Handbook contradicts both halves. Name the disagreement rather than picking a side by tone of voice. Second, feasibility favours the never-stop position: high-sensitivity TSH assays are obtainable at most laboratories and thyroid replacement is among the cheapest, most widely stocked drugs in Indian practice, so a 6-monthly test and a daily tablet are realistic at district level.

6.7 The other direction: a low TSH on lithium

Hyperthyroidism on lithium is uncommon and, on the best evidence, incidental rather than caused. One group compared adults with at least two TSH or lithium measurements over three decades against a very large unexposed population: after adjustment for age, sex and a diagnosis of diabetes, lithium use was not associated with hyperthyroidism, with an odds ratio of 1.22 and a 95% confidence interval of 0.96 to 1.55. That study operationally defined hyperthyroidism as a TSH below 0.2 mIU/l for a database search: a study convention, never the diagnostic threshold.

A low TSH follows the same logic as a high one. In every such instance the lithium continues while further measures such as T4 and T3 are arranged with the primary care physician, a repeat TSH is agreed, and a persistently low result triggers the work-up appropriate to the cause: antithyroid peroxidase antibodies for Hashimoto's disease, **thyroid-stimulating antibodies** for Graves' disease. Treatment then follows the aetiology and is unaffected by the presence of lithium. As with the initiation interval, the source directs you to agree a repeat interval without saying what it is.

6.8 Thyroid hormone augmentation in unipolar refractory depression

The same axis now becomes a treatment, and in unipolar depression the strategy is triiodothyronine augmentation. Thyroid hormone has been added to antidepressants since the 1960s and, like lithium augmentation, has faded in popularity as newer options arrived. Trials comparing thyroid with lithium augmentation directly have generally suggested thyroid is at least as effective, and the few head-to-head trials have tended to favour it. One caveat is often dropped: most of the evidence base was generated with tricyclics rather than SSRIs, though a number of trials support the strategy with SSRIs. Indian consultation-liaison teaching corroborates the indication, describing T3 supplementation as an established strategy for augmenting antidepressants where response is inadequate, without attaching a dose.

The agent is **liothyronine**. Case reports and small studies support 25 to 50 mcg of liothyronine added to various SSRIs as a way of potentiating their antidepressant effect, and Kaplan and Sadock's Comprehensive Textbook of Psychiatry gives the same range for mood-disorder augmentation: **25 to 50 mcg as a single morning dose**, timed to avoid insomnia. T3's short half-life argues for divided dosing, but single daily dosing protects adherence; the

recommendation is written for healthy adults and needs modifying in older patients and in significant medical illness.

That range is not a flat band. For augmentation in refractory depression, 25 microgram per day sits below the replacement dose for the hormone, making thyroid toxicity extremely rare, while 50 microgram per day approaches the daily physiological production range, where symptoms of thyroid excess can appear. Continuous clinical monitoring and TSH assessment are recommended while T3 is being administered, at every dose in the range and not only at the top of it. Separately, and because there are no well-demonstrated dose-response studies for T3, doses in excess of 25 microgram per day should be used with caution. The two ends differ in safety, not merely in strength, but the monitoring floor does not move with the dose.

Among the largest and most systematic trials of thyroid hormone augmentation of an SSRI to date was the **STAR*D trial**, in which patients who had failed citalopram and then bupropion or buspirone augmentation received either triiodothyronine or lithium. About 25% of those given T3 augmentation of citalopram reached a remission-level response, against about 15% of the lithium-augmented patients. Both are approximations, and the trial was open-label with patients able to decline randomisation, which makes the comparison suggestive rather than decisive.

6.9 Bipolar disorder, the dose envelope and the safety limits

Thyroid abnormalities have long been recognised as destabilising in bipolar disorder, associated with rapid cycling, mixed states, depression and refractory mania, and supplementation has been used principally in two states: rapid cycling and bipolar depression. Indian consultation-liaison teaching independently links rapid cycling and lithium refractoriness with clinical and subclinical hypothyroidism. Thyroid hormone augmentation here runs at doses of a different order from the unipolar ones.

INDICATION	AGENT AND DOSE	EVIDENCE BASE
Unipolar refractory depression, added to an antidepressant	Liothyronine 25 to 50 mcg as a single morning dose.	Case reports, small studies, one large open-label augmentation trial.
Bipolar depression, added to an existing regimen	T3 to an average dose of 90.4 mcg.	Retrospective series of 159 patients averaging 14 previous failed trials; Eighty-four percent improved on a Clinical Global Impression scale, 33% in remission.
Rapid-cycling bipolar disorder, supraphysiological T4	Levothyroxine 100 to 300 mcg per day.	Small studies from the 1990s, mostly open-label or small case series.
Bipolar disorder, T4 as adjunctive mood stabiliser	T4 in doses ranging from 150 to 500 microgram per day.	Five case studies or series totalling 30 patients, of whom 27 responded.
Added to lithium, high-dose T4 or T3	Dosed to suppress TSH below 0.1, a figure the source reports without a unit.	A 2018 randomised controlled trial in which T4, but not T3 or placebo, produced less time depressed or mixed and more time euthymic.

Two features of that table are examinable. First, the same textbook gives two ranges for **supraphysiologic doses of levothyroxine** in rapid cycling, in chapters describing overlapping but non-identical evidence sets. They must not be merged: the merged range is one no source states, and it misrepresents the top of the envelope widely. Second, the randomised evidence in bipolar disorder favours T4 over T3, the reverse of the unipolar picture.

PITFALL

Carrying the bipolar T3 average into unipolar augmentation, or a bipolar T4 dose into unipolar refractory depression. That average, roughly double the top of the unipolar range, comes from an uncontrolled retrospective series in an extraordinarily treatment-resistant population: a description of what was done, never a recommendation. The distinct T4 dose for unipolar augmentation is not established here, so the strategy is named and the dose withheld.

The apparent superiority of T3 over T4 in unipolar depression should be held loosely. The literature has not settled it: most studies used T3 and the only head-to-head comparison favoured T3, but it ran for 2 weeks, too short against T4's long half-life to establish superiority, and positive T4 studies exist. On who benefits, responders tend to be female with a history of thyroid abnormality, and one report picked out T3 responders by lower T3 levels and by carriage of a particular allele at a single nucleotide polymorphism for the **deiodinase** converting enzyme, a finding awaiting replication.

Tolerability of thyroid hormone augmentation sorts into three scopes and they must not be merged. Thyroid supplementation has generally been well tolerated in most of the studies to date,

and a transient headache, usually subsiding after one or two days, is the most commonly reported adverse effect on T3. Bound to dose: the potential side effects of high-dose thyroid supplements include osteoporosis, tachycardia, palpitations and tremor. Bound to no dose at all: T3 should be used with great caution in older adults and in patients with various medical disorders, including endocrine and cardiovascular disease, and should be avoided in pregnancy. Bound to duration: long-term thyroid hormone treatment may be associated with bone demineralisation, resulting in osteoporosis in women. That last hazard explains the duration limit: prolonged administration to a euthyroid person is not advisable, and in a depressed patient who has responded, treatment should be **tapered and stopped by 12 weeks**. Two boundaries close the topic. No data support the efficacy of thyroid hormone for depression in children. And nothing in the hypothalamic-pituitary-thyroid axis carries regulatory approval anywhere for treating depression, although thyroid hormones, chiefly T3, do appear in clinical practice guidelines for major depressive disorder: T3 as an augmenting agent in refractory depression and a possible accelerator of antidepressant response, T4 as an unsettled option in refractory depression better placed as an adjunct to mood-stabiliser treatment, particularly in rapid-cycling patients.

> Hypothyroidism on lithium is supplemented, never a reason to stop the drug: TSH and calcium every 6 months, L-thyroxine started on the TSH. In the other direction, T3 is the unipolar augmenting agent and T4 the bipolar and rapid-cycling one, and every thyroid dose you quote must carry its indication with it.

7 · Subclinical thyroid disease, autoimmunity and Hashimoto encephalopathy

This is where a thyroid report stops being reassuring. Overt hypothyroidism and overt thyrotoxicosis announce themselves, and where this chapter deals with hypothyroidism and hyperthyroidism it sets out how the function tests are read. The harder half is the patient whose free hormones are normal, whose report therefore reads as unremarkable, and in whom the thyroid is nevertheless part of the psychiatric problem. Two constructs sit in that gap. The first is **subclinical hypothyroidism**, the psychiatrically important form of subclinical thyroid disease and a family of abnormalities occurring without overt functional hypothyroidism, repeatedly implicated in depression that will not shift on adequate treatment. The second is **Hashimoto encephalopathy**, an uncommon autoimmune encephalopathy that can present to a psychiatrist as psychosis, agitation or a depressive illness, and is treatable when recognised.

AS HIGH AS 10%	9.9 mU/L	>95%	>1,000 IU/l
POPULATION BASE RATE, ANTITHYROID ANTIBODIES	UPPER BOUND, GRADE 1 SUBCLINICAL BAND	SEIZURES IN HASHIMOTO ENCEPHALOPATHY	ANTI-TPO TITRE RAISING ITS LIKELIHOOD

7.1 What "subclinical" actually names

What the word strictly means is biochemical, and Kaplan and Sadock's Comprehensive Textbook of Psychiatry subdivides subclinical thyroid disease into three grades, each with a different investigative anchor. **Grade II hypothyroidism** is an elevated thyroid stimulating hormone without any change in thyroid hormones. **Grade III** is an abnormal thyroid stimulating hormone

response to stimulation with thyrotropin releasing hormone, so it is invisible on a basal panel. **Grade IV** is the presence of antithyroid antibodies with no abnormality of the thyroid hormone system at all.

That sequence is also the order in which the grades become clinically ambiguous: by the fourth you are reading an immunological marker in a person whose thyroid is working.

The word is a poor one: it promises the absence of symptoms, does not deliver it, and is taken sometimes as a reason not to treat and sometimes as a reason not to look further. Yet middle-aged patients carrying subclinical hypothyroidism may have measurable cognitive dysfunction and nonspecific symptoms of the kind that fill psychiatric clinics, fatigue and altered mood in particular. The term describes the laboratory, not the person.

7.2 Two grading schemes, overlapping numerals

This is the most avoidable error in the topic, and purely a matter of vocabulary. The grades above are one scheme. A second labels bands of thyroid stimulating hormone concentration and also begins its numbering at a low integer. **Grade 1 subclinical hypothyroidism**, in the sense used by The Lithium Handbook, is a concentration band: thyroid stimulating hormone lying between the upper limit of the reference range, typically quoted as either 4.5 or 5.0 mU/L, and 9.9 mU/L. The source hedges that lower bound between two values rather than fixing one, which tells you how much weight a single decimal deserves.

LABEL	SCHEME	WHAT DEFINES IT	HOW IT IS FOUND
Grade II hypothyroidism	Kaplan and Sadock	Elevated thyroid stimulating hormone, thyroid hormones unchanged	Basal thyroid function panel
Grade III	Kaplan and Sadock	Abnormal thyroid stimulating hormone response to thyrotropin releasing hormone	Dynamic stimulation test only
Grade IV	Kaplan and Sadock	Antithyroid antibodies present, thyroid hormone system normal	Serology only
Grade 1 subclinical hypothyroidism	The Lithium Handbook	Thyroid stimulating hormone from the upper limit of the reference range to 9.9 mU/L	Basal thyroid stimulating hormone value

TRAP **Answering "grade two" or "grade one" without saying whose grading you are using.** The two schemes are not versions of one another and are not numbered on the same principle: one grades the depth of investigation required, the other a concentration band. A candidate who names neither produces an unmarkable answer, and one who knows only one scheme may map a numeral onto the wrong definition. Whenever you use a grade, name the scheme in the same breath.

7.3 Subclinical hypothyroidism and the depression that will not move

Grade II hypothyroidism has been associated with depressive disorders, and the association runs both ways: patients with major depressive disorder show an increased incidence of it, and some studies find these patients respond poorly to conventional treatment. The stronger formulation the standard reference offers is that Grade II hypothyroidism may itself be a risk factor for major depressive disorder rather than merely a correlate.

The phenomenology is not confined to mood. Cognitive disturbance, memory in particular, and psychotic symptoms have both been reported in Grade II hypothyroidism. The presentations attributed to it include the two symptom domains a psychiatrist is least willing to lay at an endocrine door.

The same association reappears in the antidepressant literature from the other end. There is growing evidence that patients with **mild thyroid abnormalities**, meaning a slightly elevated thyroid stimulating hormone with normal triiodothyronine and thyroxine, are less likely to respond to antidepressants given alone, but may respond to the addition of a thyroid supplement. One demographic signal travels with this: women over the age of 50 appear to show higher rates of response to thyroid augmentation, perhaps because they are the group most susceptible to hypothyroidism. The doses and the choice of agent for augmentation belong where this chapter deals with lithium, thyroid dysfunction and thyroid hormone augmentation.

> **GOOD TO REMEMBER**
>
> The practical consequence is a habit, not a drug. In any depressive episode that has failed an adequate trial of an adequately dosed antidepressant, look at the thyroid stimulating hormone value again rather than at the free hormones. A mildly raised value with normal hormones is precisely the report that gets filed as normal, and precisely the one that predicts the non-response you are trying to explain.

7.4 The same number means different things in different patients

Here the topic stops being a matter of cut-offs and becomes a matter of populations. The Lithium Handbook makes the most examinable argument in the subclinical half of this material. Its Grade 1 band, in the general population, is rarely associated with somatic, cognitive or mood effects. Read alone, that argues for leaving these patients alone. But the same passage immediately qualifies it for one group: the threshold for thyroid supplementation may be different in patients with bipolar disorder, because thyroid stimulating hormone values sitting at the upper limit of the normal range have been associated with more depressive relapse in that population. The same source states the general position elsewhere: the association between hypothyroidism and depression is not robust in the general population, but patients with bipolar disorder are unique in this regard.

Notice what that does to a reference range. The assay is not wrong: a population-derived reference interval was never designed to answer a question about relapse in a mood disorder, and using it as though it were is a category error. Lithium as a drug, its therapeutic range and its non-thyroid organ effects belong to the Mood disorders: pharmacotherapy notes.

RULE **A thyroid number is only interpretable against the population it was measured in.** The subclinical band that is clinically inert in a general medical population is not demonstrably inert in bipolar disorder, and no single threshold can be quoted as universal. Asked for a cut-off, give the band and then the population qualifier; a bare number is a wrong answer wearing a decimal point.

7.5 Should it be treated? The sources do not agree

Whether subclinical thyroid disease should be treated is a live disagreement inside the standard references and must be presented as one. Three positions are on the table.

- The endocrine account in Kaplan and Sadock's Comprehensive Textbook of Psychiatry is cautious. It concedes that there is some evidence of improvement in cognition, mood and psychosis, and then states that treatment of subclinical hypothyroidism remains controversial.
- The therapeutic-indications account in the same textbook is materially more permissive: the presence of depressive symptoms in a patient with subclinical hypothyroidism could itself warrant starting thyroxine. Same book, different chapter, different threshold for acting.
- The Lithium Handbook does not adjudicate between them but reframes the question, arguing that the threshold is population-dependent and that the unsettled debate about normal ranges is itself an argument for a flexible treatment approach rather than a fixed rule.

The third position can be read as a resolution of the first two, and the reading is productive: if the answer depends on whether the patient has bipolar disorder, a cautious general stance and a permissive stance in symptomatic patients are not contradictory. Offer that as a synthesis, not as though a single source said it, because none did.

TRAP **Answering "yes, treat" or "no, do not treat" as though the question were settled.** Candidates reproduce whichever position they read and are marked against the other, because both sit in the same standard reference and neither can be dismissed. The mark is in naming the disagreement, stating that benefit in cognition, mood and psychosis is acknowledged while the treatment question is left explicitly controversial, and then offering the population-dependent reading as a synthesis rather than a citation.

7.6 Antithyroid antibodies: what a positive result is worth

The Grade IV category is where psychiatric reasoning most often goes wrong, and the reason is a base rate. Elevations of **antithyroid autoantibodies** occur in the general population at a rate **as high as 10%**. Against that background a positive antibody in a psychiatric patient tells you very little alone.

The associations define the company this finding keeps. Elevated serum antithyroid antibodies are associated with other thyroid disorders, including Grave disease, de Quervain thyroiditis, primary hypothyroidism and colloidal goitre, and with other autoimmune disorders including diabetes mellitus, Addison disease and pernicious anaemia. A positive result is therefore a prompt to ask what else is autoimmune in this patient and this family.

The second reason for restraint is mechanistic, and applies with most force to the encephalopathy below. The role of antithyroid antibodies is itself controversial: there is no direct evidence that these autoantibodies exert direct effects on central nervous system tissue, and they may be epiphenomena of some other autoimmune process. The antibody is associated with the syndrome and used to identify it, and may have nothing to do with causing it. The third reason is evidential: these sources carry no figure connecting thyroid antibodies and treatment-resistant depression, so a positive serology has no quantified predictive value for antidepressant non-response, and the link you can defend in an examination is the one that runs through subclinical hypothyroidism above.

RULE **An antithyroid antibody is a marker, not a mechanism and not a diagnosis.** Its base rate in ordinary people is high enough that a bare positive explains nothing, and no direct action on central nervous system tissue has been demonstrated. What converts it into information is context: the titre, the thyroid hormone status, the presence of an encephalopathic syndrome, and the company of other autoimmune disease.

7.7 Hashimoto encephalopathy: the syndrome and its striking frequencies

Hashimoto encephalopathy, also described as **steroid-responsive encephalopathy associated with autoimmune thyroiditis**, is best defined as an uncommon autoimmune encephalopathy of unknown aetiology associated with high titres of serum antithyroid antibodies, usually antithyroid peroxidase with or without anti-thyroglobulin. Note that it is defined by antibody titre rather than by thyroid function. The second name is the clinically useful one, because it puts the treatment response into the label.

The syndrome is commoner in women, with a mean age of onset between **45 and 50 years**, and there may be a history of other autoimmune disease. That is the age band where a psychiatrist is otherwise thinking about perimenopausal mood change, an early neurodegenerative process, or late-onset psychosis.

The core picture is an encephalopathy with progressive cognitive decline, although the course may equally be relapsing and remitting, itself a useful discriminator against a steadily progressive degenerative illness.

- Seizures, in more than 95% of cases.
- Stroke-like episodes, in more than 65% of cases, alongside memory dysfunction.
- Visual hallucinations, frequently reported at more than **90%**, with other disorders of perception and delusions also described; agitation and restlessness, apathy and social isolation are the accompanying neuropsychiatric features, and the syndrome may present as a depressive illness.

MNEMONIC

Seizures, Strokes, Slowing, Steroid-responsive

The four features no rival explanation ties together: **seizures** in almost every reported case, **stroke**-like episodes in the majority, electroencephalographic **slowing** tracking severity, and a **steroid**-responsive course. A psychiatric presentation carrying any two deserves the antibody.

The stroke-like episodes are a feature of this disease, not a separate cerebrovascular event to be worked up as one; cerebrovascular syndromes belong to the Stroke, TBI, delirium and focal brain syndromes notes, as does the delirium workup a fluctuating encephalopathy will otherwise trigger. Where the presentation is progressive cognitive decline, the dementia syndrome and its investigation belong to the Dementias and neurocognitive disorders notes. What belongs here is that this cause is treatable.

7.8 Titre, electroencephalogram and the evidence of reversibility

Two investigations carry this diagnosis. The first is the antibody titre, where the base rate does its work. A typical psychiatric presentation such as depression, in a patient with mildly elevated antithyroid peroxidase antibodies, is unlikely to be due to Hashimoto encephalopathy. When those levels are very high, meaning **greater than 1,000 IU/l**, the likelihood of association with a neuropsychiatric Hashimoto encephalopathy is much greater. Read that as a likelihood modifier, which is what the source says, and not as a diagnostic criterion, which it does not.

The second is the electroencephalogram. Routine investigations are often normal, but patients will often show slowing on the electroencephalogram, with the degree of slowing proportional to the clinical severity of the syndrome, and some will show triphasic waves. The proportionality makes the electroencephalogram a severity measure rather than a binary test, and a way of following a treated patient when the mental state is hard to score.

QUESTION	ANTIBODIES WITH NORMAL THYROID HORMONES	HASHIMOTO ENCEPHALOPATHY
Prior probability	Common: base rate as high as 10% in the population	Uncommon
Antibody titre	Not specified by the source for this category	Mild elevation argues against it; greater than 1,000 IU/l makes it much more likely
Thyroid hormones	Normal by definition	Not stated in the source
Clinical picture	Mood, memory and occasionally psychotic symptoms	Encephalopathy with seizures, stroke-like episodes and visual hallucinations
Electroencephalogram	No signature described by the source	Slowing proportional to severity, sometimes triphasic waves
Treatment implication	Contested	Majority respond to corticos-teroid treatment

What makes the syndrome worth chasing is documented reversibility, illustrated in the standard reference by two individual cases rather than a series. In a woman of 38 with psychosis and an antithyroid peroxidase titre of 779, the pretreatment recording showed general slowing with high-voltage delta biphasic and triphasic waves at 2 to 3 Hz; after corticosteroid treatment it returned to alpha frequency with occasional posterior theta waves at 5 to 6 Hz. In a woman of 59 with rapidly progressive cognitive impairment and myoclonus, single-photon emission computed tomography showed gross global hypoperfusion in all nonoccipital regions, and her Mini-Mental State Examination score rose from 20 to 27 over the 6 weeks as she improved significantly. The instruments and their scoring belong to the Psychotherapies, clinical assessment, MSE and nosology notes; the direction and timescale belong here.

PITFALL

Treating the high-titre threshold as a rule-out. The illustrative case above carried a titre of 779, below the level at which the source says likelihood becomes much greater, and she nonetheless had a steroid-responsive encephalopathy with electroencephalographic reversal. The threshold raises probability; it does not close the door. The clinical syndrome, not the titre, puts this diagnosis on the list.

7.9 Management, and the honest limits of what the sources supply

LOCUS. The two halves are managed in different places. Subclinical thyroid disease in a mood-disorder patient is an outpatient problem, worked through over successive visits with repeat biochemistry and, where treatment starts, joint endocrine input. Suspected Hashimoto encephalopathy is not: an encephalopathic patient with seizures and stroke-like episodes needs inpatient assessment with neurology, and where consciousness or seizure control is compromised the locus is acute medical care. The governing authorities are Kaplan and Sadock's Comprehensive Textbook of Psychiatry and The Lithium Handbook, which do not agree on the first question.

FOCUS. In subclinical disease there is no acute target. The short-term target is the correct attribution of fatigue, low mood and cognitive complaint, and confirmation that the biochemical picture is persistent rather than a stray value. In the medium term it is treatment resistance itself; in the long term relapse, which is where the bipolar qualifier bites. In Hashimoto encephalopathy the acute targets are seizure control and the encephalopathy, the short-term target is response to immunosuppression, and the longer-term target is the relapsing and remitting course.

MODUS. For subclinical hypothyroidism the pharmacological mode is thyroxine, on the indication the more permissive textbook position accepts. Where a mild thyroid abnormality coexists with antidepressant non-response, the described mode is adding a thyroid supplement rather than switching; the agents and doses belong where this chapter deals with lithium, thyroid dysfunction and thyroid hormone augmentation. Psychological and social modes are unchanged by the endocrine finding and should not be paused while it is investigated. For Hashimoto encephalopathy the pharmacological mode is corticosteroid treatment: the majority of patients respond to corticosteroid treatment with complete resolution of the neuropsychiatric symptoms.

That sentence is the entirety of what these sources say about treating it. They name no agent, no dose, no route, no duration and no second-line immunotherapy. Neither should you: state the steroid responsiveness and that the regimen is set with neurology, and do not invent a milligram figure for a condition defined by its steroid responsiveness.

IN INDIA

Three practical divergences shape this topic here. First, the antithyroid peroxidase assay is widely available in private laboratories but is often reported qualitatively as positive or negative, and the high-titre threshold depends entirely on a quantitative value in stated units, so the request must specify one. Second, in a self-paying patient the serology and the neuroimaging together cost more than the consultation that generated them, so sequencing matters: in an encephalopathic patient the electroencephalogram is cheaper, more widely available, and the investigation the standard reference actually anchors. Third, iodine status and thyroid autoimmunity vary across regions, so a local reference interval is not interchangeable with a Western textbook one; where a decision hangs on a marginal thyroid stimulating hormone value, use the reporting laboratory's interval and repeat the sample before acting.

Subclinical means the hormones are normal and something else is not; a mildly raised thyroid stimulating hormone predicts antidepressant non-response, an antibody at a population base rate as high as 10% means little on its own, and a very high antithyroid peroxidase titre in an encephalopathic patient means Hashimoto encephalopathy until steroids prove otherwise.

8 · Hyperparathyroidism, hypoparathyroidism and calcium-related psychiatric symptoms

Calcium is the electrolyte psychiatry forgets to measure. The **parathyroid glands** are small accessory glands, anatomically bound to the thyroid but functionally separate from it, and they have one job: to hold serum calcium inside a narrow range so that nerve and muscle work properly. They do it by releasing **parathyroid hormone**, which raises serum calcium by stimulating osteoclasts to liberate it from bone and by increasing its absorption in the gut and the kidney. A control system whose declared purpose is to protect the excitability of the nervous system will, when it fails, fail there first. Too much calcium damps excitability towards apathy, slowed cognition and eventually obtundation; too little releases it towards spasm, seizure and confusion.

The marks sit in four places. The two directions of derangement have different signatures, one indolent and easily missed, the other neuromuscular and dramatic. Both are broadly reversible, which is what makes them worth finding. The link between biochemistry and psychiatric picture is a gradient, not a switch. And lithium sits across the whole topic, with a literature that does not agree with itself.

50%	UNDER 10%	39%	ABOUT 4%
HYPERPARATHYROIDISM, ASYMPTOMATIC	PSYCHIATRIC SYMPTOMS, HYPERPARATHYROIDISM	COGNITIVE IMPAIRMENT, HYPOPARATHYROIDISM	HYPERCALCAEMIA ON LONG-TERM LITHIUM

8.1 The set point, and why a psychiatrist should care

Most endocrine axes interest psychiatry because a hormone acts on the brain. This one is different: the hormone acts on an ion, and the ion determines how readily neurones and muscle fibres fire at all. Two consequences follow. The psychiatric picture is non-specific and syndromally promiscuous, because a drifting calcium produces whichever syndrome that brain is disposed to produce rather than a signature delusion, so the differential is wide and the diagnosis is usually made by the laboratory, not the interview. And the disturbance should track the biochemistry in time, worsening as the calcium drifts and settling as it is corrected. That temporal coupling separates a metabolic presentation from a primary illness running in parallel with an incidental abnormal result.

8.2 Hyperparathyroidism: how it is actually found

The teaching picture is florid. The reality is that hyperparathyroidism is usually diagnosed after an incidental finding of hypercalcaemia on routine blood tests, with half of all patients asymptomatic: a biochemical diagnosis handed to a clinician who was not looking for it. The commonest mechanism is increased parathyroid hormone release from a single functioning adenoma, though multiple adenomas occur as part of **Multiple Endocrine Neoplasia Syndrome**, which is why a young patient with hypercalcaemia earns wider endocrine and family enquiry than an older one.

The somatic symptoms, when they appear, are secondary to the hypercalcaemia rather than the hormone: fatigue, general malaise, proximal muscle weakness, renal colic, abdominal pain and cognitive decline. That list repays memorising precisely because it reads so easily as depression with somatic preoccupation in a tired middle-aged patient.

Background rates set how seriously to entertain it. Hypercalcaemia in the general population runs at approximately 1% to 2%, with 90% of those cases accounted for by primary hyperparathyroidism and malignancy-associated hypercalcaemia together. **Primary hyperparathyroidism** itself ranges from **0.2% to 0.8%** of the general population. Both are **2.5-fold** commoner among women and age dependent, with a 3-fold higher prevalence in patients over 80 years of age than in those aged 20 to 29. These describe the general population, not lithium-treated patients, and that distinction matters later.

8.3 The psychiatric syndrome of hypercalcaemia

The mnemonic is famous and the prevalence behind it is not what it implies. Although a range of psychiatric symptoms have been described in hyperparathyroidism, the overall prevalence is likely to be less than 10%. Hold that against the memorability of the tetrad and you have the shape of the question: the psychiatric limb is the best-known part of a syndrome in which it is the least common.

MNEMONIC

Bones, stones, moans and psychic groans

The symptomatic tetrad of hyperparathyroidism, pointing at the skeleton, the renal tract, the abdomen and the mind. Use it as a recall aid for where to look, never as a statement of frequency: the reference that prints this mnemonic is the same one that deflates the psychiatric prevalence to under a tenth of patients.

Within that minority the pattern is orderly. Affective disorders, predominantly depressive, are the most frequently recorded, with anxiety often comorbid; psychotic disorders are reported but infrequent, with persecutory and paranoid delusions predominating; cognitive changes are common, usually short-term memory loss, disorders of attention and acute confusional states. At the severe end, severe derangement of calcium metabolism may present with somnolence and coma. Symptoms correlate with both the duration and the severity of the hypercalcaemia, so a modest elevation of long standing and a sharp recent rise are different clinical problems.

PITFALL

Neuropsychiatric symptoms are often the initial presentation in older adults or those with limited cognitive reserve, which makes the commonest real-world error here not a wrong answer but an absent test. An older adult with months of low mood, apathy, poor attention and short-term memory loss is an ordinary referral, and the formulation writes itself as a depressive episode with cognitive impairment or an early neurodegenerative process. Both are likelier; only one of the three is fully reversible. The syndromal workup of dementia belongs to the Dementias and neurocognitive disorders notes and the cognitive examination to the Psychotherapies, clinical assessment, MSE and nosology notes, but the calcium is a decision you make.

8.4 Hypoparathyroidism and the neuromuscular signature

Where hyperparathyroidism is usually found by accident, hypoparathyroidism is usually made by a surgeon. The most common cause of impaired parathyroid hormone production in adults is inadvertent surgical removal during thyroid surgery, or excessive removal in the operative treatment of hyperparathyroidism itself. That single fact tells you whom to suspect and gives the history its natural question.

The clinical features follow the hypocalcaemia and are dominated by **neuromuscular excitability**: paraesthesias, muscle cramps, carpopedal spasm and facial grimacing, progressing to laryngeal spasm and convulsions. Examination may reveal **tetany**, reduced or absent deep tendon reflexes, papilloedema and prolongation of the QT interval on the electrocardiogram. Two of those are easy to lose. Papilloedema means fundoscopy, skipped far more often in a psychiatric assessment than performed. The QT prolongation matters twice over here, since it sits in the same column as the drugs a psychiatrist might reach for in a patient presenting with spasm, anxiety and paraesthesias who looks, at first pass, like panic disorder.

Note what the reference does not say. It gives this examination list without naming any eponymous elicited sign, so the two signs every candidate reaches for are not asserted here. Know

the phenomenon, latent neuromuscular irritability made manifest, rather than a name the source does not carry.

8.5 The neuropsychiatric syndrome of hypocalcaemia

The psychiatric yield here is substantially higher than on the hypercalcaemic side, which is the reverse of what most candidates expect. Delirium is now understood to be the most common neuropsychiatric manifestation. Delirium as a syndrome, its criteria, bedside detection and workup belong to the Stroke, TBI, delirium and focal brain syndromes notes; what this chapter owns is that a metabolic cause of it sits in the parathyroid axis and answers to correction of the calcium.

Beyond delirium, one large study demonstrated cognitive impairment in 39% of patients, affective or neurotic symptoms in 12%, psychotic symptoms in 11% and nonspecific affective disturbance in 21%. Those are not small numbers for an endocrine disease, and the overall yield plainly exceeds the under-a-tenth quoted on the hypercalcaemic side. The rank order differs as well: there affective disorders are the most frequently recorded, whereas in this study cognitive impairment is the largest single group. No per-symptom breakdown is published for the hypercalcaemic side, so the two cannot be compared any more finely than that. As on the other side, severity relates directly to the degree of hypocalcaemia and appropriate normalisation resolves the symptoms.

GOOD TO REMEMBER

Persistent psychosis has been noted in hypoparathyroidism where associated **hypomagnesaemia** was not addressed. If you have normalised the calcium and the psychosis has not moved, check the magnesium before you raise the antipsychotic. A failure to respond is information about what has been left uncorrected before it is information about treatment resistance.

8.6 A gradient, not a threshold

This is where the chapter deliberately stops short. Both directions carry an explicit severity statement in the standard reference: the more severe the hypercalcaemia, the more severe the psychiatric disturbance, and on the other side symptom severity relates directly to the degree of hypocalcaemia. Both are made without a single calcium figure attached. There is no value in these sources at which psychiatric symptoms are said to emerge, none at which they worsen, and none at which somnolence and coma supervene.

RULE **Psychiatric severity tracks the calcium, and the calcium is what reverses it, in both directions.** This one principle does the work of a dozen memorised facts: it tells you what to measure, it warns you that a mild derangement with a florid psychiatric picture needs a second explanation, and it predicts that correction is the treatment. What it does not do is give you a number, so teach the gradient qualitatively. A remembered threshold, printed with a citation borrowed from the sentence next door, is how a false figure enters a textbook and survives in it.

Two related absences are worth naming, because both are slots where a plausible number would sit unchallenged. No reference range for total or ionised calcium is stated in these sources, and

neither is the albumin-correction formula, notwithstanding that **albumin-adjusted calcium** is an outcome measure in the lithium literature below.

One further discrimination protects against a real error of attribution. It is the calcium that drives the psychiatric picture, not the parathyroid hormone. When hormone levels are mildly raised in the setting of normocalcaemia, most commonly on the basis of **vitamin D insufficiency**, psychiatric disturbance is uncommon.

TRAP **Treating a raised parathyroid hormone as the psychiatric abnormality.** Candidates make this error because the disease is named after the gland and the hormone is what was measured, so it feels like the causal variable. It is not the one the neuropsychiatric syndrome tracks. The source describes mildly raised hormone with normal calcium, the picture produced by vitamin D insufficiency, as increasingly prevalent in developed countries, and psychiatric disturbance as uncommon there. A stem offering a raised hormone, a normal calcium and a psychiatric syndrome is usually offering a syndrome that belongs to something else.

8.7 Lithium and calcium: a literature that disagrees with itself

Lithium is why this topic appears in a psychiatry paper at all, and the mechanism is better established than the magnitude. Sufficient human and animal data indicate that lithium exposure can induce hypercalcaemia through increased gastrointestinal and renal calcium reabsorption, a direct impact on **calcium sensing**, and secondary or independent effects on parathyroid gland function. Altering the calcium-sensing mechanism can itself produce hyperparathyroidism, so the drug reproduces the endocrine disease rather than merely nudging the ion. Lithium as a drug, its pharmacokinetics, therapeutic range and non-thyroid organ-system adverse effects belong to the Mood disorders: pharmacotherapy notes; taught here is the calcium lens only.

The size of the effect is another matter. Four standard sources address it and they do not line up.

SOURCE	QUANTITY MEASURED	FINDING
2012 systematic review and meta-analysis, n = 699 across 4 cohort, 14 case-control, 36 case-report and 6 cross-sectional studies	Mean change in total calcium and in parathyroid hormone	Total calcium raised by **0.09 mmol/l** (95% CI 0.02 to 0.17, p = 0.009); parathyroid hormone raised by 7.32 pg/ml (95% CI 3.42 to 11.23, p < 0.0001)
Oxford University Hospitals NHS Trust laboratory database, n = 1916	Risk of raised albumin-adjusted calcium	No increased lithium-related risk, OR 1.08, 95% CI 0.88 to 1.34, p = 0.46
Maudsley prescribing guidance	Proportion of lithium-treated patients affected	Hyperparathyroidism causes hypercalcaemia in about 4% of patients
Geriatric psychiatry guidance, meta-analysis of 385 studies	Magnitude of the rise in blood calcium	Blood calcium increased by **10%**

The honest reading is that the direction of effect is agreed and almost nothing else is. Two of these are effect sizes, one an incidence and one a proportional rise; the largest single dataset finds no signal; the two study counts cannot be reconciled from what is printed. The handbook itself

names the discrepancy between the meta-analysis and the larger database study and leaves it open, and a revision chapter that closes it by choosing a favourite is doing what the primary literature declined to do.

TRAP **Reading a proportion of patients and a proportional rise as the same quantity because both wear a percent sign.** About 4% of lithium-treated patients developing hypercalcaemia and a 10% increase in blood calcium answer entirely different questions, and side by side in a revision note they read either as a contradiction needing resolution or, worse, as corroboration. They are neither. Any stem offering a single percentage for lithium and calcium is under-specified until you know whether it counts patients or measures calcium. The same discipline applies to the general-population rates given earlier: those describe everybody, not lithium-treated patients, and the source is explicit that they must not be attached to a lithium patient. No source here states which way that error would run, so do not reason from a direction; counsel a lithium patient from the lithium-treated figure.

One apparent contradiction resolves cleanly. The lithium handbook states plainly that lithium does not cause osteoporosis, while the prescribing guidance lists osteoporosis among the consequences of chronically raised serum calcium. These are compatible if the routes differ: what is denied is a direct skeletal action of the drug, what is asserted is a consequence of sustained hypercalcaemia however caused. State them together without that distinction and you have written a contradiction into your own notes.

8.8 Management: finding it, correcting it, monitoring for it

LOCUS. Almost all of this work is outpatient: incidental hypercalcaemia, surveillance on long-term lithium, and follow-up after thyroid surgery are clinic problems. Inpatient care is needed where the psychiatric syndrome itself is unsafe, meaning the confusional and somnolent presentations of severe hypercalcaemia and the delirious and psychotic presentations of hypocalcaemia, and the medical or endocrine ward is usually the right ward, because the treatment is correction of biochemistry. Emergency presentation belongs to the neuromuscular end of hypocalcaemia, where laryngeal spasm and convulsions are why this is not a problem to observe over a weekend.

FOCUS. Acutely the target is the calcium, not the symptom, and the severity relationship is the reason. In the short term it is the cause: the adenoma on the high side, replacement after surgical loss of the glands on the low side. In the medium term it is the resolution that should follow correction, whose failure to arrive is itself a finding, and where magnesium is checked. In the long term it is surveillance, because the consequences of sustained hypercalcaemia accumulate quietly.

- Renal stones
- Osteoporosis
- Dyspepsia
- Hypertension
- Renal impairment

MODUS. Definitive treatment on the hyperparathyroid side is surgical, and the reference is explicit that symptoms generally respond to appropriate treatment such as parathyroidectomy. Where surgery is not the route, **calcimimetic agents** offer a non-surgical option, which is the relevant sentence where lithium is the driver and an operation is a large ask. Psychiatric treatment is adjunctive and symptomatic, and its principle is restraint: a syndrome expected to resolve on correction of the calcium does not need a durable psychotropic regimen built around it. The psychological and social modes are largely explanatory, since a patient told that months of low mood and forgetfulness had a correctable cause needs that explained if the episode is not carried forward as a psychiatric history.

Monitoring is where the guidance is most usable and least unanimous. Both prescribing sources agree on the principle: routine serum calcium monitoring will identify the adverse effect, with **ionised calcium** and parathyroid hormone used for confirmation, and calcium monitored on long-term lithium. On timing they diverge. The geriatric guidance directs clinicians to measure serum calcium before lithium is initiated and annually thereafter, while the lithium handbook recommends a shorter interval this chapter does not print, not being among the values it can vouch for. Take the baseline, follow whichever interval your local protocol specifies, and treat a calcium before a first lithium prescription as non-negotiable, because without it a later raised value cannot be attributed to anything.

IN INDIA

The divergence here is test ordering rather than pharmacology. Lithium monitoring in most Indian settings reliably captures the lithium level, renal function and thyroid function, and calcium is the parameter most often dropped, partly because the psychiatric consequence is unfamiliar and partly because it is another line on a bill paid out of pocket. The consequence is specific: the incidental route by which hyperparathyroidism is usually diagnosed does not operate unless somebody ordered a calcium, so a patient whose only regular biochemistry is a lithium clinic panel has no incidental route at all. Total calcium with albumin is cheap and widely available; the confirmatory ionised calcium and parathyroid hormone assays sit in tertiary centres and larger private laboratories and often need a send-out, so plan that step when you order the screen. One caution: the source used here attaches the rising prevalence of vitamin D insufficiency to developed countries and offers no Indian figure, so none is manufactured.

Psychiatric severity tracks the serum calcium in both directions and reverses when it is corrected, so a calcium belongs in the workup of any unexplained affective, psychotic or cognitive presentation and in every lithium patient at baseline; no source here gives a threshold at which symptoms begin, and none should be invented.

Metabolic and nutritional disease

9 · Wilson's disease: copper metabolism, genetics and the neuropsychiatric presentation

This is the diagnosis that arrives in a psychiatric clinic wearing someone else's clothes. A young adult with irritability, emotional lability and a falling academic trajectory; or a first psychotic episode; or a tremor dramatic enough to be called functional. The referral is written for a mood disorder, the examination stops above the neck, and a treatable single-gene disorder of copper metabolism runs on unrecognised.

What follows is the neuropsychiatric presentation of Wilson's disease and the genetics behind it. Confirmatory biochemistry and the chelating drugs are treated where this chapter deals with the investigation and management of Wilson's disease.

9.1 What Wilson described, and what has been added since

The condition carries the name of Kinnear Wilson, who first recognised the syndrome in 1912 and included in that original description a high rate of psychiatric abnormality. Psychiatric involvement was in the founding account, alongside the movement disorder and the cirrhosis, not a later refinement.

Cummings identified the role of copper in pathogenesis in 1948. Schenberg and Gitlin described the deficiency of ceruloplasmin in 1952, establishing a test still in use. Bearn established the recessive nature of the disorder definitively in 1960. The molecular basis emerged with linkage to the esterase D locus on chromosome 13 in 1985, and then, in 1993, with discovery of the causative mutations in ATP7B. Note that order: ceruloplasmin deficiency was described four decades before the gene was found, which is why older texts treat it as the primary lesion.

9.2 The gene, the protein and the transport failure

Wilson's disease is an autosomal recessive disorder caused by mutations in a copper transporting ATPase encoded by the **ATP7B** gene on chromosome 13q14.3. Kaufman's Clinical Neurology for Psychiatrists localises the gene to chromosome 13, coding for a protein necessary for copper transport, without naming ATP7B; the sub-band locus comes from the more precise source.

The protein matters more than the locus, because copper metabolism fails at a single step. The **Wilson protein** is an ATPase involved in copper transport whose physiological job is to maintain copper homeostasis by facilitating excretion of copper from hepatocytes into bile. Biliary excretion is the only meaningful route for surplus copper, and that route fails, so that copper accumulates in liver, brain and other tissues, producing hepatic failure together with neurological and psychiatric signs and symptoms, with the basal ganglia appearing particularly affected within the brain.

The order is the whole of the mechanism question. The gene defect leads to accumulation of copper in the liver through impaired copper excretion and impaired binding of copper to ceruloplasmin. Catabolism of the resulting ceruloplasmin produces the low plasma ceruloplasmin

level, with a rise in free copper, and free copper then accumulates in the brain and generates the neurological and neuropsychiatric manifestations.

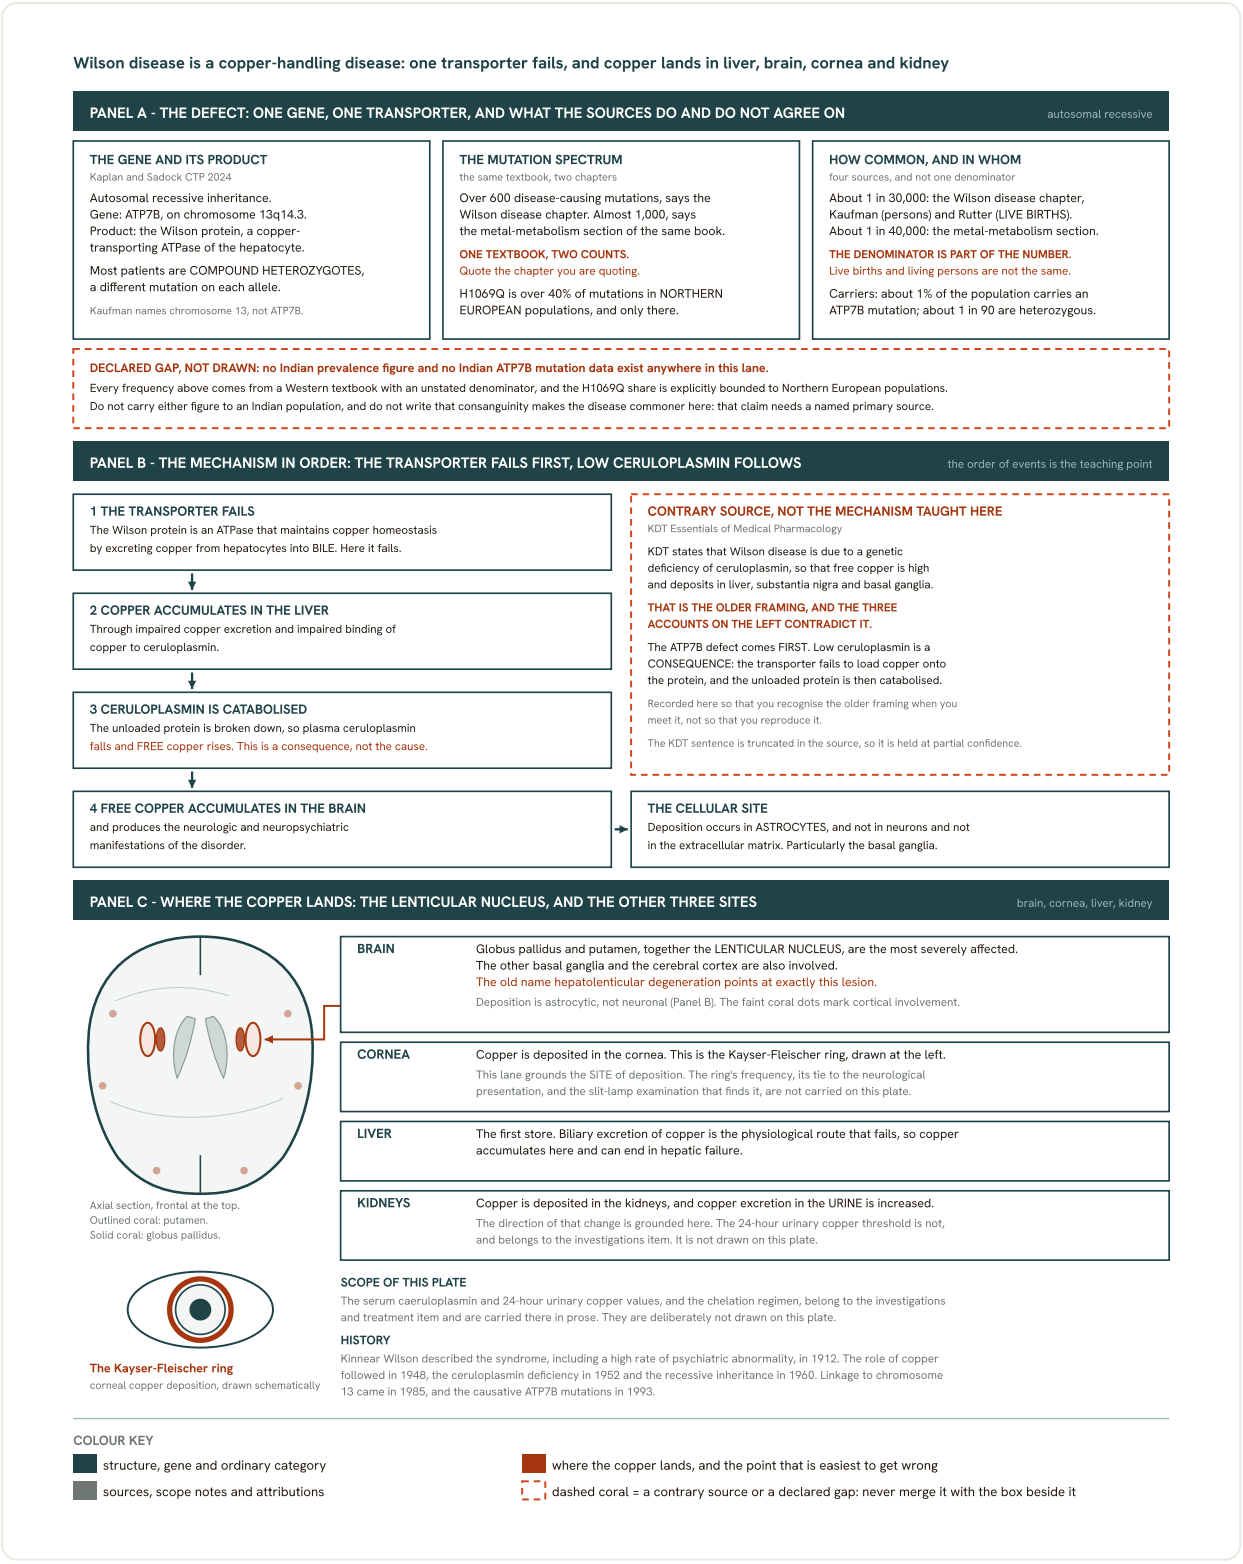


Fig 33.4 Wilson disease as a copper-handling disease: the ATP7B defect, where the copper lands, and the lenticular nucleus. The order of events is the teaching point: the transporter fails first, copper accumulates in the liver because biliary excretion and loading onto ceruloplasmin both fail, the unloaded protein is catabolised, and only then does plasma ceruloplasmin fall and free copper rise. Low ceruloplasmin is a consequence, not the primary lesion, and the older pharmacology framing that reverses this is drawn as a contrary source. Cerebral deposition is astrocytic and falls hardest on the globus pallidus and putamen, which is what the name hepatolenticular degeneration points at. The frequency figures disagree across sources and denominators and are shown unmerged; no Indian figure exists in this lane and none is drawn. (Kaplan and Sadock CTP 2024; Kaufman's Clinical Neurology for Psychiatrists; Companion to Psychiatric Studies; Rutter's Child and Adolescent Psychiatry 6e; KDT Essentials of Medical Pharmacology.)

RULE **Low ceruloplasmin is a consequence of the ATP7B gene defect, not its cause.** Everything downstream follows from surplus free copper, which is why the disease is treated by removing copper and not by replacing ceruloplasmin.

TRAP **Writing that Wilson's disease is due to a genetic deficiency of ceruloplasmin.** This inversion is not invented by candidates, it is printed. KDT's Essentials of Medical Pharmacology states that the disease is due to genetic deficiency of ceruloplasmin, the protein that normally binds and disposes of copper, in whose absence plasma free copper is high and is deposited in liver, substantia nigra and basal ganglia, causing local degeneration. Kaplan and Sadock, the Companion to Psychiatric Studies and Brazis all place the primary lesion in the copper-transporting ATPase, with the ceruloplasmin fall arising secondarily. Name the ATPase account as the current one and the ceruloplasmin-deficiency account as the older framing; drop neither.

9.3 Where the copper lands, and why the disease has two names

Deposition is neither diffuse nor random. The Companion to Psychiatric Studies states that mutation in the copper-transporting P-type ATPase coded on chromosome 13 results in excessive copper deposition in brain, cornea, liver and kidneys, with increased urinary copper excretion, and that the globus pallidus and putamen, together the **lenticular nucleus**, are most severely affected, though the other basal ganglia and the cerebral cortex are also involved. That licenses the older name hepatolenticular degeneration, explains why the syndrome is extrapyramidal rather than pyramidal, and accounts for the cognitive and behavioural layer over the movement disorder.

One cellular detail is a favourite discriminator: copper deposition occurs in **astrocytes**, not in neurons and not in the extracellular matrix, and is particularly evident in the basal ganglia. The figure sets the transport defect, the deposition sites and the ocular sign side by side.

9.4 Mutation spectrum, carrier burden and a frequency figure that will not settle

This is not a one-mutation disease. The dedicated chapter of Kaplan and Sadock records that over 600 disease-causing mutations have been identified in ATP7B; a few are common, with H1069Q accounting for more than 40% of mutations in Northern European populations, while most are rare, so that most individuals with the disease are **compound heterozygotes**, carrying a different mutation on each allele. A negative test for one common variant therefore excludes very little.

Carrier burden is high for a rare recessive disorder. One chapter puts carriage at about 1% of the population, the dedicated Wilson chapter at **approximately 1 in 90 individuals**: close, not identical, from different chapters of one book. The mutation count disagrees more sharply, the first chapter saying almost 1,000 mutations have been identified against the dedicated chapter's count of disease-causing mutations. Identified and pathogenic variants are different denominators, so the two may not be counting the same thing. Quote each in its own wording, and never present the larger number as a count of pathogenic variants.

Disease frequency is the conflict most likely to be asked, and the denominators differ as well as the numerators.

SOURCE	STATED FREQUENCY	DENOMINATOR, EXACTLY AS GIVEN
Kaplan and Sadock, disorders of metal metabolism section	About 1 in 40,000.	General population.
Kaplan and Sadock, dedicated Wilson disease chapter	About 1 in 30,000, with approximately 1 in 90 heterozygous carriers.	Worldwide.
Kaufman's Clinical Neurology for Psychiatrists	1 per 30,000 persons.	Persons.
Rutter's Child and Adolescent Psychiatry	About 1 in 30,000.	Live births.

Three readings give **about 1 in 30,000** and one the higher denominator, which sits in the same textbook as one of the lower ones. Print the commoner one, note the variant, and carry the denominator with the number: a rate per live birth and a rate per general population are not interchangeable in a disease with substantial childhood mortality. Kaufman gives the reason it is quoted at all: despite its infrequent occurrence neurologists routinely test for Wilson disease in adolescents and young adults who develop dysarthria, tremor, parkinsonism, dystonia, atypical psychosis, dementia or cirrhosis. The dementia syndrome and its workup belong to the Dementias and neurocognitive disorders notes; a rare disease is worth testing for when the test is cheap and the treatment works.

IN INDIA

Every frequency figure here comes from Western textbooks, and the only mutation-frequency figure is bounded explicitly to Northern European populations. No Indian prevalence estimate and no Indian ATP7B mutation spectrum can be quoted from these sources. Do not assert that Wilson's disease is commoner in India on account of consanguinity: plausible and widely repeated, it needs a named Indian series before it belongs in an answer. Nor carry the Northern European mutation figure into an Indian population. What travels is the method: in most Indian settings the diagnosis is secured on examination, the slit lamp and biochemistry rather than sequencing, and the highest-yield single step needs only an ophthalmology colleague.

9.5 When it presents, and through which door

The presenting-system split is the most quotable statistic in the topic. About 40% of patients first present with liver abnormalities, about 40% with neurological abnormalities, and up to 20% with primarily psychiatric abnormalities. One patient in five arrives first in a psychiatric clinic, at an age when psychiatric illness is expected and metabolic disease is not.

40%	40%	UP TO 20%	17 YEARS
LIVER FIRST	NEUROLOGICAL FIRST	PSYCHIATRIC FIRST	MEAN AGE OF ONSET

Onset is usually in the second or third decade with a mean age of onset at 17 years, atypical late-onset cases are reported in individuals over the age of 60, and without treatment death may occur

within 5 years, sooner where there is central nervous system disease. Age does not exclude the diagnosis.

Natural history by decade adds the clause that matters most. Hepatic disease is the commonest presentation and occurs anywhere between the first and fourth decades; about 50% of patients are symptomatic by the age of 15; and up to two-thirds present with neurological disease in the second or third decade without clinical evidence of liver disease. Normal liver function tests and an unremarkable abdomen do not exclude the diagnosis in a young patient with a movement disorder or behavioural change, and missed cases are missed on exactly that reasoning.

9.6 The neurological syndromes and the tremor that looks functional

The neurological picture is multifocal, following the deposition pattern. Neurological signs correlate strongly with brain and cerebrospinal fluid copper levels, reflect the multiple brain regions involved, and include dysdiadochokinesis, dysarthria, tremor, dystonia, rigidity, choreoathetosis, bradykinesia, masked facies and micrographia. Seizures occur in 6% of patients. Two phenotypes are conventionally separated, and the separation is prognostically load-bearing.

- The **dystonic form** is the commonest presenting neurological syndrome, with dysarthria, dysphagia, drooling and a rigid open mouth.
- The **pseudosclerotic form** presents with incoordination, clumsiness, unsteadiness of gait or dysarthria.
- As the illness progresses patients may show features of both, with rigidity, tremor and choreic, athetoid or dystonic movements.

The face is distinctive across a consulting room. Brazis records excessive grinning, a sustained open-mouth smile sometimes called the vacuous smile, and a fixed forced smile produced by sustained spasm of the risorius and zygomaticus, termed **risus sardonicus**, alongside slow saccadic eye movements and, rarely, ophthalmoplegia.

The most examined sign is the tremor. The tremor of Wilson's disease, the so-called **wing beating** tremor, may be interpreted as a functional movement disorder. It is characteristically absent at rest and develops after a short period of arm extension; the arms beat in a wide violent arc, and the tremor may be altered by the position of the arms. To a clinician who expects organic signs to be constant, that looks manufactured. It is not. Brazis gives the manoeuvre: a large-amplitude wing-beating tremor may be demonstrated with the shoulders abducted to 90 degrees. Kaufman notes that this hallmark occurs in only about one-third of cases and consists of rhythmic, coarse, up-and-down arm movements based at the shoulders. Keep those figures apart: one is a manoeuvre angle, the other how often the sign is present.

MNEMONIC

Two forms, one arc

Dystonic: mouth, speech and swallow, the commonest presenting syndrome. **Pseudosclerotic**: gait, coordination and speech, the cerebellar-looking one. **One arc**: the wing-beating tremor, absent at rest, appearing on sustained arm extension, and present in a minority of cases, so that its absence excludes nothing.

9.7 The psychiatric presentation, and a disagreement about psychosis

Two-thirds of patients may have psychiatric symptoms by first presentation, and 10% to 20% seek psychiatric treatment before the diagnosis is made, a stage at which the diagnosis is often missed; eventually psychiatric manifestations affect nearly all patients, and the psychiatric profile generally resembles that of Huntington disease. The other chapter reports a range: estimates range from 30% to 100% of symptomatic patients experiencing psychiatric symptoms at some point, up to two-thirds present with psychiatric symptoms, and one-third will have received psychiatric treatment before the diagnosis is made. Print that range as a range: collapsing it to a midpoint destroys what it carries.

What the symptoms are is more consistent across sources than how often they occur. The commonest are personality change, often described as irritability, aggressiveness and emotional lability, and depression ranging from subsyndromal mood change to major depression, at times with suicidality; anxiety disorders, psychotic disorders and individual psychotic symptoms, and catatonia all appear at higher rates than in the general population. The referral pattern follows: the commonest reasons for psychiatric referral are behavioural and personality change, with disinhibition or bizarre or impulsive behaviour in about a quarter of patients and depression in about one-fifth, and personality and behavioural change, but not depression, correlate with the degree of neurological impairment.

That correlation is stated directly elsewhere: psychiatric symptom severity correlates more strongly with neurological deficits than with liver disease, indicating that the psychiatric phenomena arise from brain damage rather than from a psychological reaction to the illness. Cognition follows the same substrate. It has not been well studied; a relatively mild degree of impairment akin to that in other basal ganglia diseases is thought to be present in about 25% of patients, worsening over time and in step with accumulating neurological symptoms. Where there is neurological involvement, the pattern is subcortical, with frontal executive deficits that correlate with the extent of cerebral involvement. Keep the source's own hedge on that figure.

TRAP **Answering that psychosis is a characteristic presentation of Wilson's disease, or that it is rare, as though either were settled.** One chapter of Kaplan and Sadock states that about 10% of patients present with psychosis and are frequently misdiagnosed with primary schizophrenia or bipolar disorder. Another chapter of the same textbook states that although early reports suggested psychotic presentation was a feature of the disease, more recent investigations have shown psychotic symptoms to be relatively rare. This is an intra-textbook contradiction on a load-bearing point. The defensible position is the one both support: personality and behavioural change are the commonest psychiatric presentations in every source, and the frequency of frank psychosis is disputed inside one standard reference.

9.8 The Kayser-Fleischer ring: site, frequency and the word pathognomonic

The **Kayser-Fleischer ring** is the single most important clinical diagnostic sign, consisting of copper deposits in the outer rim of the cornea that appear brown or grey-green on slit-lamp examination. Brazis gives the precise layer: deposition of copper in the **Descemet membrane**, accompanied in some patients by sunflower cataracts. Kaufman adds the caveats: it is most obvious at the superior and inferior corneal margins and usually obscures the fine structure of the iris; a slit lamp is required to see it in its early stages; its size and density correlate with disease duration where the brain is affected; and it dissolves with successful treatment.

Four standard texts describe one sign in four subtly different ways.

SOURCE	SITE AS STATED	FREQUENCY	SPECIFICITY AS STATED
Kaplan and Sadock, disorders of metal metabolism section	Outer rim of the cornea.	About half of patients with hepatic presentation; almost all with a neurological or psychiatric presentation.	Single most important clinical diagnostic sign; pathognomonicity not claimed.
Kaplan and Sadock, physical examination chapter	Limbus of the cornea, brownish-green discoloration.	Not stated.	Stated to indicate Wilson disease sensitively and specifically.
Kaufman's Clinical Neurology for Psychiatrists	Periphery of the cornea, bilateral.	About 70% of cases where the disease affects only the liver; almost 100% where it affects the brain.	Explicitly not pathognomonic, since it also develops in primary biliary cirrhosis and several other liver diseases.
Rutter's Child and Adolescent Psychiatry	On the iris.	Not stated.	Described as pathognomonic.

The iris reading is probably an artefact of appearance: Kayser-Fleischer rings lie in front of the peripheral iris and obscure it. On frequency the sources agree where it matters most to a psychiatrist, because **almost 100%** of patients with brain involvement carry the ring; they disagree only on the hepatic-only denominator, where the populations described are not identical.

TRAP **Writing that the Kayser-Fleischer ring is pathognomonic of Wilson's disease, or that it sits on the iris.** Both errors are printed in respectable texts. The safe answer names the cornea, specifically Descemet membrane, calls the ring the single most important clinical sign, and states on Kaufman's authority that it is not pathognomonic.

GOOD TO REMEMBER

Kayser-Fleischer rings are not reliably visible to the naked eye early on, so a normal-looking iris on ward examination is not a negative result: if Wilson's disease is on the differential, the request is for a slit-lamp examination by an ophthalmologist. They also resolve with treatment, so their absence in a patient already on therapy carries no diagnostic weight.

9.9 Recognising it in a psychiatric clinic, and what changes on Monday

The systemic reach of copper toxicity is wider than the classical triad suggests. Liver dysfunction, the commonest childhood presentation, may appear as mild transaminase elevation, chronic active hepatitis, cirrhosis, or acute and fulminant liver failure; copper toxicity may also produce cardiac arrhythmias, renal disease, rhabdomyolysis, joint disease, anaemia, and thyroid and parathyroid dysfunction. Note the endocrine entry: a movement disorder with a thyroid or parathyroid abnormality in a young patient is two findings that one diagnosis explains.

LOCUS. The neuropsychiatric presentation of Wilson's disease is recognised in the outpatient psychiatric clinic, the consultation-liaison referral, and occasionally the emergency department; confirmation and treatment move into shared care with hepatology and neurology. The psychiatrist's task is the same at each: suspect the diagnosis and arrange the examination that supports it, not complete the workup alone.

FOCUS. Immediately the target is recognition, because the importance of recognising the early psychiatric presentations lies in the proven benefits of early treatment intervention. In the short term, the confirmatory workup and transfer into shared care. Longer term the psychiatric symptoms need treating in their own right; they do not remit the moment a diagnosis is made.

MODUS. No treatment guideline for this disease is quoted here; the sources behind the above are Kaplan and Sadock, Kaufman, Brazis and the Companion to Psychiatric Studies. Clinically, extend the examination: eye movements, tone, gait, speech, the face, and a request for slit-lamp examination. Biochemically, serum ceruloplasmin, urinary copper, the confirmatory sequence and chelation are set out where this chapter deals with the investigation and management of Wilson's disease. Psychiatrically, symptomatic treatment continues alongside; where antipsychotics are involved, their own movement-disorder effects are the subject of the Schizophrenia: management, antipsychotics and adverse effects notes.

PITFALL

Attributing new extrapyramidal signs in a young psychiatric patient to their medication. Diagnostic delay is exacerbated when extrapyramidal movements are attributed to psychotropic medication or interpreted as functional. Both are reasonable on their face and both self-sealing: once the tremor is drug-induced, nobody examines the cornea. Signs that preceded the drug, that persist after it is withdrawn, that come with dysarthria, dysphagia or drooling, or that are unusually dramatic in a young patient, deserve a slit lamp before they deserve an anticholinergic.

Wilson's disease is an autosomal recessive failure of the ATP7B copper-transporting ATPase: biliary excretion fails, ceruloplasmin falls as a consequence and not a cause, and free copper accumulates in astrocytes of the lenticular nucleus and in Descemet membrane of the cornea. Up to a fifth of patients present psychiatrically, almost all of those with brain involvement carry a Kayser-Fleischer ring, and normal liver function tests exclude nothing.

10 · Wilson's disease: investigations (ceruloplasmin, urinary copper) and management (chelation)

No single test makes this diagnosis. Wilson's disease is confirmed on a panel of imperfect measures, and every element has a documented failure rate an examiner expects you to name. The copper-handling defect, the genetics, the Kayser-Fleischer ring and the presentation are set out where this chapter deals with Wilson's disease as a disorder of copper metabolism; what follows is the work-up and the treatment.

The work is worth it: this is one of very few neurodegenerative diseases in which copper accumulation, the fundamental pathogenic process, is itself amenable to treatment, so early consideration of the diagnosis, including in older patients, is essential. The standard references also contradict themselves in four places here; in each, the safe answer names the disagreement rather than picking a side.

95%	10%	UP TO 50%	1 MONTH
CERULOPLASMIN BELOW NORMAL	NORMAL CERULOPLASMIN, HIGH URINE COPPER	NEUROLOGICAL WORSENING, PENICILLAMINE	WINDOW FOR BEST OUTCOME

10.1 Serum ceruloplasmin: the first test, and the ways it misleads

The screening test is **serum ceruloplasmin**, the copper-carrying protein of plasma, which in this illness is low to absent. One mechanism explains both copper tests: the defective transporter fails to load copper onto the apoprotein, the unloaded protein is cleared, and the uncarried copper is excreted, so one lesion lowers the serum protein and raises the urine copper. It is no marker of brain disease, since even where the disease does not affect the brain, ceruloplasmin concentrations fall to very low levels.

Kaplan and Sadock's Comprehensive Textbook of Psychiatry puts the level below normal in 95% of patients, which sounds decisive until the same sentence lists what else lowers it. It is low in newborns and in **20% of heterozygous carriers**, and it may be misleadingly normal in the

setting of infection or inflammation. The mechanism matters: ceruloplasmin is an acute-phase reactant, so sepsis or hepatitis can lift a diagnostic value into the reference range. There the source redirects you, since the absence of serum ceruloplasmin oxidase activity is a useful test. Measuring **ceruloplasmin oxidase activity** rather than protein mass sidesteps it: what is missing is the functional copper-loaded enzyme.

One absence must be stated rather than filled: no source used here gives a numeric cut-off or any units, saying only below normal, low, or low to absent. Take any printable threshold from a named primary hepatology guideline.

RULE **A normal serum ceruloplasmin does not exclude Wilson's disease.** Both chapters of the textbook agree on this while disagreeing on the arithmetic behind it. A patient with a normal result and a fitting picture still needs the urinary copper, the slit-lamp and, if the question stays open, confirmatory testing.

10.2 Urinary copper, and what else raises it

The second core investigation is the **24-hour urinary copper**, and the textbook position is strong: an elevated 24-hour urinary copper concentration is generally considered diagnostic of the disease. Kaufman's Clinical Neurology for Psychiatrists explains why, and why clinicians avoid it, calling it a more cumbersome test, in which the absence of ceruloplasmin greatly increases urinary copper excretion. Cumbersome is the word: an incomplete collection reads falsely low in a patient who has the disease.

Here too the sources give direction and withhold the number: none states a threshold, its units, or a provocation protocol, and that matters more than the ceruloplasmin absence, because the value means nothing without a laboratory reference. The result is less specific than that wording suggests, since urine copper excretion may be elevated for reasons unrelated to Wilson's disease:

- Metal chelation therapy already in progress.
- Acute liver failure, or cirrhosis with cholestasis.
- Nephrotic syndrome.

Hepatic copper has its own mimics: it may be elevated in cholestatic liver disease, where ceruloplasmin is normal or elevated, and in idiopathic childhood cirrhosis, where ceruloplasmin is elevated. The ceruloplasmin direction separates them, which is why the panel is read together. The mimic to memorise is **aceruloplasminaemia**, which produces neurological symptoms but raises liver iron rather than liver copper: what discriminates is the metal, not the missing protein.

PITFALL

Sending the 24-hour urinary copper after chelation has started. It happens whenever a confident diagnosis leads to treatment while confirmatory tests are still in the system, and it destroys the result: chelation is itself a cause of raised urinary copper, so the sample no longer separates disease from drug.

10.3 Confirming the diagnosis: liver biopsy against gene sequencing

Where the copper studies leave the question open, two confirmatory routes exist and the textbook cannot decide between them. Its Wilson disease chapter holds that liver biopsy is the gold standard for diagnostic confirmation, and may show increased hepatic copper concentration with evidence of hepatitis or cirrhosis, though no threshold and no units accompany it anywhere in these sources. Its metal-metabolism chapter measures it differently, giving liver biopsy with staining for total hepatic copper a sensitivity of over 80%, with false negatives arising from sampling error: copper is unevenly distributed through a cirrhotic liver, so a core can miss the loaded parenchyma.

The molecular route is the **ATP7B** gene. Sequencing is confirmatory when disease-causing mutations are identified, but a negative test, absent supportive laboratory data, does not exclude the disease, because the mutations are heterogeneous. The metal-metabolism chapter goes further, calling genetic testing of that gene the gold standard, while noting that it requires direct sequencing of the entire gene and may return variants of unclear significance.

TRAP **Naming one gold standard for Wilson's disease and defending it.** "Gold standard" is the phrase an examiner reaches for, and the same textbook answers it two ways in two chapters. The mark is in the reasoning: say that histological and molecular confirmation are both called the reference test, that each has a stated weakness (sampling error for the biopsy, uninterpretable variants for the sequence), and that they are complementary.

10.4 Neuroimaging and the pathological substrate

Imaging supports the diagnosis and tracks treatment but does not screen: MRI may help diagnose the illness, but characteristic changes often do not emerge until irreparable damage has occurred. A normal scan argues for the copper studies, not against the diagnosis.

The signature to memorise is reduced T1 signal and increased T2 signal in the basal ganglia, thalamus and brainstem. Kaplan and Sadock add generalised atrophy with increased T2-weighted signal in the basal ganglia secondary to copper deposition, sometimes in patients with no clinical evidence of brain disease: the brain may be involved radiologically before clinically. Atrophy may also involve the cerebral cortex, particularly the frontal lobe, and the brainstem and cerebellum; behind the images lies cavitary degeneration of the basal ganglia, with neuronal loss, gliosis and high tissue copper.

Localization in Clinical Neurology is more precise: striatal lesions are the most characteristic lesions found in the brain in patients with neurological symptoms, and MRI demonstrates bilateral symmetric hyperintensities in the putamen, thalami, brainstem and posterior limb of the internal capsule, symmetry discriminating against most other striatal lesions in a young adult. The lesions can be more diffuse, involving the pons, midbrain and dentate nucleus, and less frequently the corpus callosum. A published axial T2 image of an adolescent shows the narrower pattern, bilateral hyperintensity of the thalamus, caudate and putamen.

The eponymous sign needs care: the sources describe it differently and merging them yields a description neither supports. Kaplan and Sadock name the **midbrain panda sign**, common in

the disease but nonspecific, from hyperintense signal around the red nucleus and putamen, while Localization in Clinical Neurology attaches double panda sign to the putaminal, thalamic and brainstem hyperintensity just described. Quote the sign with its source, and carry the qualifier Kaplan and Sadock attach to theirs: suggestive, not diagnostic. Localization in Clinical Neurology states its sign flatly, with no comment on specificity, so the hedge belongs to the first source and should be cited to it.

10.5 Assembling a diagnosis from imperfect tests

Because no test stands alone, someone had to formalise the combination. The **Leipzig criteria** were developed by the European Association for the Study of the Liver, incorporating clinical and laboratory findings, in recognition of how variable both are in this disease. That is all these sources state: no items, no point weights, no total and no interpretation, so a scoring table from memory would be ungrounded. Take the arithmetic from the primary guideline.

TEST	WHAT THE STANDARD SOURCES STATE	THE STATED LIMITATION
Serum ceruloplasmin	Low to absent, even where the brain is unaffected.	Low in newborns and carriers; may be normal in inflammation.
24-hour urinary copper	An elevated concentration is generally considered diagnostic.	Cumbersome. Raised by chelation, acute liver failure, cholestatic cirrhosis, nephrosis.
Liver biopsy	Increased hepatic copper with hepatitis or cirrhosis. Sensitivity over 80%.	False negatives from sampling error.
ATP7B sequencing	Confirmatory when disease-causing mutations are identified.	A negative result does not exclude it. Variants of unclear significance.
MRI	Reduced T1, increased T2 in basal ganglia, thalamus, brainstem.	Characteristic changes often do not emerge until irreparable damage has occurred.

TRAP **Quoting a single sensitivity figure for serum ceruloplasmin.** The same textbook gives two, in two chapters, and they cannot be reconciled. One puts the level below normal in 95% of patients. The other states that low serum ceruloplasmin has high sensitivity, over 80%, with low specificity, and that a normal ceruloplasmin may occur in 10% of patients in the presence of high urinary copper excretion, a sensitivity that is not arithmetically consistent with its own false-negative figure. The two chapters therefore imply different false-negative rates. Give the direction, name both as variants, and land on the consequence both support: you cannot rule the disease out on this test. The passage adds that low serum copper or free copper and high urine copper can be seen but are less reliable.

10.6 The aims of treatment and the agents that serve them

Treatment rests on three aims and every drug maps onto one: to remove accumulated copper, to prevent re-accumulation through maintenance treatment, and to reduce its absorption.

MNEMONIC

Strip, Sustain, Starve

Strip out what is stored with a chelator; **sustain** the gain with maintenance; **starve** the body of new copper by blocking absorption. The agent names the phase your patient is in.

The mainstay chelator is **d-penicillamine**. The Companion to Psychiatric Studies agrees that chelating agents are effective but carry a risk of significant side effects, penicillamine being currently the most widely used, with trientine and zinc acetate as alternatives. KDT Essentials of Medical Pharmacology describes it as dimethyl cysteine, a degradation product of penicillin, used for Wilson's disease from 1956, which selectively chelates Cu, Hg, Pb and Zn, plus the examination favourite: the d-isomer is used therapeutically because the l-isomer and the racemate produce optic neuritis and are more toxic. In practice, penicillamine combined with pyridoxine has been the mainstay of initial treatment, though no source used here states either dose, so neither is printed.

Trientine, or triethylenetetramine dihydrochloride, is an acceptable alternative for patients with Wilson's disease intolerant of penicillamine, and is effective orally. For that indication, copper chelation in penicillamine-intolerant Wilson's disease, Goodman and Gilman gives maxima of **2 g per day** in adults and 1.5 g per day in children, in two to four divided portions on an empty stomach. Those are maxima, not starting doses, and no titration schedule appears in the source.

Zinc acetate serves the third aim by a mechanism that fixes its place: it induces metallothionein in the gastrointestinal epithelium and thereby inhibits copper absorption, and has been used as preventative treatment in presymptomatic patients and for maintenance. The trapped copper is shed with the enterocyte, so zinc blocks rather than de-coppers, which is why it belongs to prevention and maintenance. No zinc dose appears in any source used here. The repertoire is stated once: treatment involves either removing copper by chelation, using D-penicillamine, trientine and tetrathiomolybdate, or blocking copper absorption with zinc salts. Tetrathiomolybdate is named there and nowhere else, with no dose.

TRAP **Writing that trientine is the safer chelator for neurological Wilson's disease.** The sources disagree about which agent carries the risk. One attributes it to penicillamine, whose combination with pyridoxine may lead to a worsening of neurologic symptoms in **up to 50% of patients, which has led to the recommended use of less neurotoxic copper chelators such as trientine**. The other gives it to trientine, which can paradoxically increase neurologic and psychiatric symptoms acutely as hepatic copper is mobilised and deposited in the brain, giving penicillamine a propensity to induce hypersensitivity reactions instead. That mechanism is generic to de-coppering and fits either drug. Name early deterioration on starting chelation as a hazard of de-coppering itself, say the sources differ on which agent carries it, and attribute the up-to-half figure to penicillamine, the agent named by the source stating it.

10.7 Managing the patient: LOCUS, FOCUS, MODUS

LOCUS. Diagnosis and maintenance are outpatient work shared between psychiatry, neurology and hepatology, and the psychiatrist often sees the patient first. Admission is for initiating

chelation in significant neurological disease, where early deterioration is a hazard, and for decompensated liver disease; transplant-centre care belongs to fulminant hepatic failure. The guidance base is thin: the European Association for the Study of the Liver stands behind the diagnostic criteria, but the treatment recommendations come from standard texts, so local hepatology protocols govern the regimen.

FOCUS. Acutely, the copper load and the behaviour that brought the patient in; short term, neurological deterioration as chelation begins; medium term, adherence, where most avoidable harm is done because the patient feels better and stops; long term, re-accumulation, liver monitoring and residual deficit.

MODUS. Drug treatment is the sequence above. Diet is where the sources conflict and both positions must be carried: one chapter holds that restriction of high-copper foods, namely shellfish, liver, chocolate, nuts, beans and mushrooms, is encouraged; the other that except in areas with high levels of copper in the drinking water, dietary restrictions are inadequate to prevent copper overload. Neither says which it means, so if you counsel restriction, say explicitly that it is adjunctive, never a substitute for chelation. Procedurally, orthotopic liver transplantation has proven successful in end-stage liver disease and restores normal copper excretion, though it is limited by lifelong immunosuppression and procedural complexity.

Psychiatric treatment has almost no trial evidence behind it, and saying so is examinable. There are no controlled trials of symptomatic treatment for the psychiatric manifestations, and anecdotal evidence suggests particular vulnerability to the motor side effects of antipsychotics; the prudent course is standard psychiatric treatment, pharmacological and psychotherapeutic, at lower doses, increased slowly, avoiding high-potency antipsychotics. Antipsychotic receptor pharmacology and drug selection belong to the Schizophrenia: management, antipsychotics and adverse effects notes; added here is the reason for the caution, a basal ganglia already loaded with copper and degenerating. No milligram figure attaches to "lower doses" in any source here.

IN INDIA

The copper-restricted food list is written for a Western plate. Nuts, beans and mushrooms are staples in most Indian households and pulses a principal protein source in largely vegetarian diets, so handing it over unmodified either goes unfollowed or does real nutritional harm; counsel specifically, with a dietitian where available. The timed collection is harder where a patient travels a long distance and must return a complete sample, and a partial one reads falsely low. Cost and stock shape the choice of chelator more than any textbook comparison: penicillamine is the cheapest and most widely stocked, and the alternatives are not uniformly available outside larger centres.

10.8 Timing, recovery and what to tell the patient

The strongest argument for psychiatric vigilance is one sentence. Treatment outcomes are best when treatment begins within 1 month after symptom onset; the average symptom duration before correct diagnosis and treatment, however, is about 1 year. The gap is where the

preventable disability is made, much of that year in psychiatric services, because behavioural change and mood disturbance often precede the movement disorder.

Recovery is real but slow, and patients should be told so before they judge the drug a failure. Neurological and psychiatric symptoms improve with copper chelation over 1 to 2 years, with regression of MRI lesions on serial imaging. The counselling figure is a year or two, not weeks, and imaging gives an objective marker of response to a patient who cannot yet feel one.

Outcome differs by phenotype, and the series must be quoted carefully. The best outcomes are seen in patients with the pseudosclerotic form (82% symptom-free), with worse outcomes in patients with chorea (75% of patients), parkinsonism (63%) and dystonia (53%). The source does not repeat "symptom-free" after the first figure, so the denominator of the last three is not explicit; they read as a descending series of symptom-free rates, but nothing says so, and relabelling them would be invention with a citation attached. Teach the ordering: the pseudosclerotic form does best, dystonia worst.

The sources extend early consideration to older patients, and the teaching that this is a disease of the young is what produces the delay. And treatment is lifelong, because copper re-accumulates when the drug stops, so the fully recovered patient is the one at risk of stopping.

> A normal ceruloplasmin does not exclude Wilson's disease, and an elevated 24-hour urinary copper is uninterpretable once chelation has begun: send both before the first dose of anything, in any young patient with psychiatric symptoms plus a movement disorder or unexplained liver disease.

11 · Porphyria and the adult-onset inherited metabolic disorders

These are the diagnoses a psychiatrist is structurally placed to miss. An inherited metabolic disorder declaring itself in infancy belongs to the paediatrician; these wait until adolescence or adult life and open with a syndrome psychiatry already names, so the psychiatrist sees them first, the label is defensible the day it is made, and it closes the case. The porphyrias are episodic disorders of heme synthesis in which neurotoxic precursors accumulate under metabolic stress; the storage and remethylation disorders are progressive.

1 IN 20,000	1 IN 500	1.5 PER 100,000	25% TO 40%
FREQUENCY OF AIP	AIP IN 4,000 PSYCHIATRIC INPATIENTS	FREQUENCY OF MLD	PSYCHOSIS, ADOLESCENT AND ADULT NPC

11.1 Acute intermittent porphyria: the defect and who gets it

The porphyrias are disorders in which defects in heme biosynthesis cause excessive urinary secretion of porphyrins and their precursors, several pathway intermediates being neurotoxic. Most porphyrias are dominantly inherited with incomplete penetrance, so the family history is often blank in a carrier.

Acute intermittent porphyria is the one that matters to psychiatry: autosomal dominant and incompletely penetrant, arising from defects in **porphobilinogen deaminase**, the enzyme

converting porphobilinogen to hydroxymethylbilane, and occurring in **1 in 20,000** individuals.

The block accumulates the precursors porphobilinogen and aminolevulinic acid, both thought to be neurotoxic. The deficit is latent: it becomes apparent in situations requiring an increase in heme synthesis, when demand pulls the pathway forward and precursors pile up behind the block. Hence who presents, since the disorder presents most commonly in women of child-bearing age, in whom one such demand recurs monthly.

- Fasting.
- Menstruation.
- Intercurrent medical illness.
- Cytochrome P450 inducers, including alcohol, estrogens, barbiturates and sulfonamides.

In 4,000 psychiatric inpatients the disorder was significantly over-represented at 1 in 500 against a community rate of 1 in 100,000.

TRAP **Quoting "1 in 500" as the prevalence of acute intermittent porphyria.** The population is part of the number: it describes psychiatric inpatients, not any general or outpatient population. A subtler version sits alongside it: the standard reference gives the disorder frequency as 1 in 20,000 and, in the same passage, a community rate of 1 in 100,000. The two do not agree and no source reconciles them. Quote either with its population attached; never average them or invent a reconciliation.

11.2 The attack: a triad, and the two diagnoses it gets instead

The **classical triad** during an episode is abdominal pain, psychiatric disturbance and peripheral neuropathy. The neuropathy is what candidates get wrong: mostly motor and often mimics Guillain-Barre syndrome, not the length-dependent sensory neuropathy met in diabetes or alcohol use.

MNEMONIC

Pain, Psyche, Peripheral nerve

Abdominal **pain**, disturbance of the **psyche**, a mostly motor **peripheral nerve** lesion. Two of the three sit outside psychiatry, and on any visit only one limb may show, since psychiatric symptoms alone may be the single presenting feature. That is why the diagnosis survives several admissions.

Among clinically symptomatic cases, psychiatric disturbance occurs in **up to half of all cases**, and half of those are psychotic episodes, with depression, anxiety and delirium also appearing as the main presenting symptom. The denominator is symptomatic cases, not all gene carriers. Delirium as a syndrome belongs to the Stroke, TBI, delirium and focal brain syndromes notes.

Two misdiagnoses recur. The intermittent character of the attacks invites a misdiagnosis of schizophrenia, and abdominal pain during attacks invites a diagnosis of somatoform disorder. The second is more damaging, because a somatoform formulation discourages further physical investigation, as with the functional movement disorder label misapplied in Wilson's disease, set out where this chapter deals with Wilson's disease and its neuropsychiatric presentation. Beyond

the attacks, the more chronic porphyrias carry elevated rates of anxiety and depression. The periodic madness of King George III has in recent decades been attributed to the disorder, which has also been implicated in illnesses affecting the work of Van Gogh and Soren Kierkegaard; hedged attributions, colour in an answer rather than fact.

11.3 Why the nervous system suffers, and how firmly we can say it

How precursor accumulation produces neuropsychiatric disturbance is unclear, and candidate explanations have included oxidative stress, vascular change and demyelination. Keep that hedge; presenting one hypothesis as established loses credit.

Aminolevulinic acid is structurally similar to gamma-aminobutyric acid, which may impair GABA release from the synapses of inhibitory neurons. Failed inhibition is what agitation, psychosis, seizure risk and autonomic instability look like together. Alongside it, reduced heme-dependent enzymes are proposed to increase serotonin turnover and reduce nitric oxide activity, the second a route to the vascular limb.

11.4 Confirming it: the specimen is half the test

Diagnosis rests on elevated urinary aminolevulinic acid and porphobilinogen together with red cell porphobilinogen deaminase activity. Gross elevations often turn the urine amber or purple in direct sunlight, the classic image and the clue to handling: if light changes the analyte, light destroys the specimen. Hence a 24-hour urine in light-protected containers, correctly preserved.

No cut-off, units or reference band is given for any of the three analytes. Do not reconstruct one; leave interpretation to the reporting laboratory.

Hereditary coproporphyria is the near neighbour: a dominantly inherited deficiency of coproporphyrinogen oxidase that can present identically, with the addition of cutaneous findings: bullous lesions, skin fragility and scarring. The neuropsychiatric picture gives nothing, so the skin discriminates and the laboratory follows it: coproporphyria calls for urinary coproporphyrin levels.

11.5 Managing the acute attack and prescribing around it

LOCUS. An acute attack is emergency inpatient medical work: a motor neuropathy mimicking Guillain-Barre carries a respiratory risk, and the treatments that reverse the biochemistry are intravenous. The psychiatrist's locus is then consultation-liaison: securing the admission, advising on psychiatric symptoms, auditing the drug chart. The independent psychiatric locus is outpatient, in preventing the next attack.

FOCUS. First, correct identification and avoidance of precipitants where possible, the only intervention that changes the trajectory rather than the episode. Acute targets are the biochemical block, the contributing illness and the agitation or psychosis; long term, a regimen that holds the psychiatric disorder without provoking attacks.

MODUS. No dedicated guideline for the psychiatric care of porphyria is available here; what follows rests on Kaplan and Sadock's Comprehensive Textbook of Psychiatry. Attack treatment includes reversal of contributing illnesses, intravenous hydration, carbohydrate loading to inhibit

heme synthesis, and hematin or heme arginate to provide negative feedback to the heme synthetic pathway. The last two signal that heme is abundant, suppressing enzymes upstream of the block so precursors stop accumulating. No dose, route, frequency or duration appears for any of them, and an invented figure here would be one for the definitive treatment of a life-threatening attack.

Psychotropic selection is where the psychiatrist owns the decision. The principle is judicious use of medication that will not worsen the biochemical deficit, and specific agents are named: chlorpromazine and droperidol for psychosis, fluoxetine for depression, lithium for mania, and lorazepam, triazolam and temazepam for anxiolysis and sedation. Lithium belongs to the Mood disorders: pharmacotherapy notes; antipsychotic pharmacology and drug selection to the Schizophrenia: management, antipsychotics and adverse effects notes. One prohibition stands against that list: like the barbiturates, chloral hydrate may precipitate attacks and is contraindicated in porphyria.

PITFALL

Treating that short list of named safe drugs as a complete safety classification, and reading absence from it as evidence of danger. The source does not classify most of the drugs a psychiatrist reaches for, including the anticonvulsant mood stabilisers and the second-generation antipsychotics, and silence about an agent is a verdict in neither direction. Check the specific drug against a maintained porphyria drug-safety database before prescribing.

11.6 Metachromatic leukodystrophy: psychosis with a white matter signature

Metachromatic leukodystrophy is an autosomal recessive, incompletely penetrant deficiency of the lysosomal enzyme arylsulfatase A, occurring in 1.5 per 100,000 individuals and first described in 1933. Inheritance changes here: these disorders are recessive, so consanguinity and affected siblings become relevant. **Arylsulfatase A** hydrolyses sulfatides, including sulfate-containing lipids of the central nervous system; sulfatide accumulates in brain, peripheral nerve, kidney and gallbladder, particularly in myelinated structures, seen as metachromatic granules on histology with widespread myelin loss. A disease of myelin presents as a disconnection syndrome, not a focal deficit.

In younger patients seizures and motor symptoms predominate; in adults psychiatric manifestations and dementia frequently precede motor symptoms, particularly with the I179S mutation, and severity is inversely correlated with residual enzyme activity. More residual enzyme means later onset and a more psychiatric, less motor picture. The adult form appears to cleave into two phenotypes.

- A predominantly motor, cerebello-pyramidal presentation.
- A predominantly psychiatric presentation.

Up to half of patients with onset between 10 and 30 years of age present with psychotic symptoms, including auditory hallucinations, systematised delusions, formal thought disorder, catatonic posturing and inappropriate affect. That is a complete description of schizophrenia in a

patient whose neurological examination will initially be normal. The figure is bound to that onset band.

Time supplies the clues. Seizures, chorea or dystonia supervene, and dementia is an inexorably progressive component presenting first as attentional disturbance, reduced processing speed and executive impairment, which may present as a phenocopy of frontotemporal dementia, or, where disinhibition is prominent, as a phenocopy of hypomania. A first manic episode dominated by disinhibition, in a young adult slowly losing executive function, should trigger this thought; the dementia syndrome itself belongs to the Dementias and neurocognitive disorders notes. Diagnosis is by reduced enzyme activity in leukocytes or skin fibroblasts, MRI supporting rather than confirming: it generally shows periventricular white matter change sparing the subcortical U-fibers, often with a frontotemporal preponderance.

11.7 Niemann-Pick type C: the eye sign that changes the diagnosis

Niemann-Pick disease type C is an autosomal recessive neurovisceral lipid storage disorder with a frequency of 1 in 100,000 live births, in which **95%** of patients have aberrations on NPC1 at 18q11-12 and 5% on NPC2 at 14q24.3. One discrimination is examined: it is biochemically and phenotypically distinct from types A and B, which result from lysosomal sphingomyelinase deficiency.

NPC1 and NPC2 handle the cyclical movement of sterols within cells, and their impairment causes late endosomal accumulation of cholesterol and gangliosides. Excess GM2 and GM3 gangliosides alter neuronal ultrastructure, and disrupted cholesterol metabolism alters tau and beta-amyloid processing, with Alzheimer-like neurofibrillary tangles, startling in a young patient. The regional pattern explains the sequence: axonal structures are affected early, then cerebellar Purkinje cells, basal ganglia and thalamic neurons, with hippocampal and neocortical regions later, the most fulminant loss being cerebellar and producing early cerebellar signs.

Clinically the disorder may present in infancy, adolescence or adulthood, but its core features are dementia, dysarthria, ataxia, vertical supranuclear ophthalmoplegia and hepatosplenomegaly, with seizures, dysphagia and pyramidal signs as it progresses. Psychosis occurs in 25% to 40% of adolescent and adult-onset cases and may precede motor and cognitive disturbance by many years. The band is bound to adolescent and adult-onset cases, and the lead time makes this a psychiatric problem: there are years in which the patient is simply a person with a psychotic illness.

RULE **Vertical gaze palsy plus psychosis means Niemann-Pick type C until proven otherwise.** The co-occurrence of vertical gaze palsy and psychosis should prompt the clinician to consider this disorder, so that appropriate treatment can be instituted. It is the highest-yield rule in this group, because **vertical supranuclear ophthalmoplegia** is easy to miss and easy to elicit. A routine examination follows a finger horizontally, which will not find it: ask for voluntary downgaze and look for a saccadic abnormality rather than gross restriction.

Confirmation runs through cell biology: a low esterification rate of exogenous cholesterol in fibroblasts, or filipin staining for lysosomal accumulation of free cholesterol, with oxidised

plasma cholesterol and direct NPC1 sequencing used more recently. **Filipin staining** is the classic confirmatory test. Imaging often shows diffuse cerebral or cerebellar atrophy, striatal, hippocampal and thalamic atrophy and callosal thinning, none of it specific alone. One published case shows what confirmation looks like: a 26-year-old man presented with movement disturbance and dysarthria after a 10-year history of treatment-resistant psychosis, with gait ataxia, jerky saccades and grossly impaired downgaze. Filipin staining showed 60% to 70% of cells with perinuclear staining against a normal figure of less than 5%, and the cholesterol esterification rate was mildly abnormal at 2.9 pmol/hr/mg on a scale where less than 2 is abnormal, 2 to 3 equivocal and greater than 3 normal, with NPC1 compound heterozygosity for G992R and R1186H. Treat that esterification band as one laboratory's reported figure inside a published case, not a general threshold.

Treatment is why the rule exists. Psychosis may require higher antipsychotic doses and a combination of antipsychotics and mood stabilisers. That direction runs opposite to the antipsychotic advice in Wilson's disease, set out where this chapter deals with Wilson's disease investigations and management. Beyond symptomatic treatment, substrate reduction therapy with the imino sugar miglustat, which crosses the blood-brain barrier and reduces ganglioside accumulation, has shown promise in reversing motor and cognitive deficits and may reduce or prevent psychiatric disturbance, with trials under way of intrathecal and intravenous cyclodextrin and acetyl-L-leucine. No dose, frequency or titration is given for any of these.

11.8 The homocystinurias: one analyte up, the other in either direction

Disorders raising blood and urinary homocysteine are grouped as the homocystinurias. The classic form is autosomal recessive, caused by a defect in the chromosome 21 gene coding for cystathionine beta-synthase, the enzyme converting homocysteine and serine into cystathionine, leading to accumulation of both homocysteine and methionine. **Cystathionine beta-synthase** sits on the transsulfuration exit, so methionine rises alongside homocysteine.

Presenting features are usually developmental delay or intellectual disability with optic lens dislocation within the first 10 years; lens dislocation is seen in almost 100% of patients older than 10, with cataract, glaucoma and optic nerve atrophy as further complications, and intellectual disability in about half of patients, declining slowly. Patients may appear Marfanoid, with arachnodactyly, kyphoscoliosis and pectus excavatum, and clotting abnormalities with increased platelet adhesiveness cause occlusion and thromboembolism with very high mortality. Intellectual disability and the thrombotic tendency discriminate from Marfan syndrome, and the thrombosis is what kills.

Early reports linked homocystinuria with psychosis, but large case series report personality disorder, behavioural disturbance such as aggression, depression and obsessive-compulsive disorder as the most common findings, with aggression and behavioural disturbance more common in those with intellectual disability. Answering "psychosis" reproduces the impression those series overturned. Diagnosis rests on elevated plasma methionine with elevated plasma and urine homocysteine; treatment works when early: methionine restriction and cysteine supplementation will generally prevent long-term sequelae if the disorder is diagnosed at birth,

and oral pyridoxine, which remethylates homocysteine to methionine, reduces both analytes in about half of patients, with folic acid and vitamin B12 possibly of further benefit. No dose is stated, and the quantities used run far above ordinary supplement doses.

The second mechanism group inverts one of the two analytes. Other autosomal recessive homocystinurias arise from defects in an alternative remethylation pathway, in which homocysteine is catalysed to methionine by either methionine synthetase or methylenetetrahydrofolate reductase; an inherited defect in methylcobalamin, the cofactor for methionine synthetase, also produces homocystinuria together with methylmalonic aciduria. These deficiencies give high homocysteine but, in contrast to cystathionine synthase deficiency, low methionine.

FEATURE	CLASSICAL HOMOCYSTINURIA	REMETHYLATION DEFECTS
Enzyme or cofactor	Cystathionine beta-synthase, chromosome 21.	Methionine synthetase, MTHFR, or methylcobalamin.
Homocysteine	Elevated, plasma and urine.	High.
Methionine	Elevated in plasma.	Low.
Extra marker	None distinctive.	Methylmalonic aciduria with a methylcobalamin defect.
Psychiatric signature	Personality disorder, behavioural disturbance, depression, OCD.	Psychosis with the less severe MTHFR form.

TRAP **Assuming every homocystinuria raises methionine.** Homocysteine is high in all of them, so it cannot separate the mechanism groups. Methionine discriminates: high in the classical transsulfuration block, low where the defect lies in remethylating homocysteine back to methionine. Candidates lose the mark by reasoning from the disorder's name, which points only at homocysteine. Reason from the pathway: block the exit and substrate dams up behind it with the product it would have become; block the recycling arm and that product cannot be made at all.

Methylenetetrahydrofolate reductase is the one member with a population-level psychiatric footprint. The less severe form of the deficiency has been strongly associated with psychosis in adolescents and adults, possibly through alterations to dopamine synthesis and myelination, and the C677T variant, which causes a 35% reduction in enzyme function, is an independent risk factor for the development of psychosis. Read that **35%** correctly: it is the reduction in enzyme function attributable to the variant, not an effect size for psychosis risk.

The clinical course and MRI findings in the methionine synthetase and methylcobalamin disorders can resemble the childhood leukodystrophies, and neuropsychiatric symptoms figure most prominently among presenting symptoms in methylcobalamin deficiency, with delays of up to 13 years between onset of neuropsychiatric symptoms and the eventual diagnosis. A delay of **up to 13 years** is a career-length misdiagnosis in a treatable disorder. This is an inherited cofactor defect, not the acquired dietary deficiency set out where this chapter deals with vitamin B12 and folate deficiency.

IN INDIA

Access to the confirming test diverges here, not access to the drug. Every disorder in this section is confirmed by an assay a district or medical college laboratory does not perform: red cell porphobilinogen deaminase, arylsulfatase A in leukocytes or cultured fibroblasts, filipin staining, plasma methionine and homocysteine, gene sequencing. These are send-away tests routed to a handful of metabolic centres, with transport, cold chain and turnaround working against the referring psychiatrist, and largely paid out of pocket. Free bedside discriminators therefore carry disproportionate weight, and the porphyria specimen instruction becomes a live failure point rather than a formality, since a light-protected, correctly preserved 24-hour collection is hard to obtain reliably in a busy general hospital.

11.9 Pattern recognition: when to stop and think metabolic

The standard reference maps each storage and metabolic disorder to its characteristic psychiatric presentation, and psychosis and mood disorder recur across the storage disorder rows. They are not the whole group's vocabulary: the same table gives Wilson's disease personality change, rarely psychosis, and the large case series in classical homocystinuria report personality disorder, behavioural disturbance, depression and obsessive-compulsive disorder as the most common findings instead. The accompanying sign, never the psychiatric presentation alone, makes the diagnosis.

DISORDER	PSYCHIATRIC SIGNATURE	WHAT DISCRIMINATES IT
Acute intermittent porphyria	Psychotic episodes, depression, anxiety, delirium; episodic.	Abdominal pain, mostly motor neuropathy; urinary aminolevulinic acid and porphobilinogen.
Hereditary coproporphyria	Identical to acute intermittent porphyria.	Bullous lesions, skin fragility, scarring.
Metachromatic leukodystrophy	Psychosis.	White matter change sparing U-fibers; enzyme activity in leukocytes or fibroblasts.
Niemann-Pick type C	Mood disorder, psychosis.	Vertical gaze palsy, ataxia, hepatosplenomegaly; filipin staining.
Classical homocystinuria	Personality disorder, behavioural disturbance, depression, OCD.	Lens dislocation, Marfanoid habitus, thromboembolism; high plasma methionine.
GM2 gangliosidosis and adrenoleukodystrophy	Psychosis and mood disorder in both.	Named in the source summary only; not developed further here.

Wilson's disease belongs in the same reflex and is set out in full where this chapter deals with its copper metabolism, genetics and neuropsychiatric presentation, and separately with its investigation and management. Triggers for pausing a routine formulation: a first psychotic or manic episode in a young adult, treatment-resistant from the outset; any psychiatric syndrome with a movement disorder, gait abnormality, neuropathy or eye sign; unexplained abdominal

pain running with episodic disturbance; cognitive decline that does not recover between episodes; consanguinity or an affected sibling.

GOOD TO REMEMBER

A two-minute addition to the examination of any young adult with a first psychotic episode. Test voluntary vertical eye movements, particularly downgaze. Watch the patient walk. Look at the light-exposed skin. Ask about abdominal pain and consanguinity. None of it costs anything.

The metabolic diagnosis is made by the company the psychosis keeps: abdominal pain with a motor neuropathy means acute porphyria, vertical gaze palsy means Niemann-Pick type C, white matter change sparing the U-fibers means metachromatic leukodystrophy, and lens dislocation with a Marfanoid habitus means homocystinuria, in which high methionine marks the classical form and low methionine the remethylation defects.

12 · Wernicke encephalopathy and the Korsakoff boundary

This is the one condition in metabolic psychiatry where the correct response to uncertainty is to treat, not to test. A malnourished drinker is admitted confused. The house officer orders a computed tomography scan, waits for a normal report before settling on a diagnosis, and writes thiamine up for the morning drug round. Ordering the scan is reasonable; waiting on it is not, and the morning drug round is not defensible at all in an emergency this fragment treats as one. Both failures are the same failure: investigation and diagnostic uncertainty are allowed to set the hour at which treatment starts, and that deferral is the standard route to a permanent amnestic syndrome.

The chronic syndrome's phenomenology and prognosis belong to the Stroke, TBI, delirium and focal brain syndromes notes; the beri-beri syndromes belong where this chapter deals with pellagra and the other vitamin deficiency syndromes.

12.1 The disease, the vitamin and the three ways deficiency happens

Wernicke's encephalopathy is an acute and reversible neuropsychiatric condition resulting from thiamine deficiency, occurring mainly but not exclusively in people with alcohol use disorders who are malnourished, and also in those with malabsorption or with repeated vomiting and diarrhoea. Reversible is the source's own word, and the emergency logic rests on it.

Thiamine is an essential co-factor for many enzymes of the glycolytic and pentose phosphate pathways, so tissues wholly dependent on aerobic glucose metabolism, holding no reserve, decompensate first, which is why an encephalopathy rather than a myopathy presents. Deficiency arises by exactly three routes: low thiamine intake, impaired thiamine absorption, or impaired thiamine utilization. The first is the drinker whose diet has been displaced by alcohol, the second malabsorption and the gut complications of drinking, the third inherited susceptibility.

Kaplan and Sadock's Comprehensive Textbook of Psychiatry states that some people are at higher risk than others for this condition because of a genetically influenced transketolase

deficiency. Two patients drinking comparably need not both develop the syndrome: intake and absorption are necessary, not sufficient.

12.2 The classical triad, and how seldom it is complete

The syndrome is taught through its **classical triad**. Saunders and colleagues in Addiction medicine describe one or more symptoms or signs from the classical triad: oculomotor abnormalities, meaning nystagmus and ophthalmoplegia with paralysis or weakness most commonly affecting lateral or upward gaze and often with diplopia, reflecting the external recti as in a sixth nerve palsy; cerebellar dysfunction, meaning ataxia; and confusion of recent onset, with disorientation, inattention and poor responsiveness. The threshold is one or more of the three, not all three.

It has to be, because the complete triad is uncommon. The full triad is seen in only **16 to 20%** of cases, and the diagnosis is often missed. Autopsy data make the point from the other direction: Wernicke's pathology is 20 times more common than the classical triad, and is found in 12.5% of people who misuse alcohol and in 0.4 to 2.8% of the general population. Those two percentages are not interchangeable and neither should be quoted without its population attached. The cost of missing it: Wernicke's encephalopathy constitutes a medical emergency with an estimated mortality of 10 to 20% if untreated, and is easily reversible with the timely administration of adequate doses of parenteral thiamine.

16-20%	12.5%	10-20%	48-72 H
HAVE THE FULL TRIAD	AUTOPSY PATHOLOGY, ALCOHOL MISUSE	MORTALITY IF UNTREATED	WINDOW BEFORE IRREVERSIBILITY

12.3 Where the lesion sits, and what imaging can and cannot settle

The lesion is periventricular and paramedian. Tasman's Psychiatry records that the condition is characterised pathologically by haemorrhage in the **mammillary bodies** and the dorsomedial nuclei of the thalamus, with patients presenting clinically with delirium, with or without nystagmus and ataxia, and left with an amnesia once the delirium clears. Kaplan and Sadock name pathological alterations in the mammillary bodies and thalamus as the substrate of that amnesia; imaging of the combined syndrome adds a third structure: volume loss of the mammillary bodies, medial thalamus and periaqueductal grey matter. Delirium as a syndrome belongs to the Stroke, TBI, delirium and focal brain syndromes notes.

Imaging follows the anatomy, and its role is confirmatory. Where Wernicke's encephalopathy is clinically suspected, treatment with high-dose intravenous thiamine should commence before any investigations are carried out. A normal computed tomography scan of the brain does not rule out the diagnosis. Magnetic resonance imaging supports it by showing structural change in the region of the hippocampus, the mammillary bodies, the mammillothalamic tract and the thalamus.

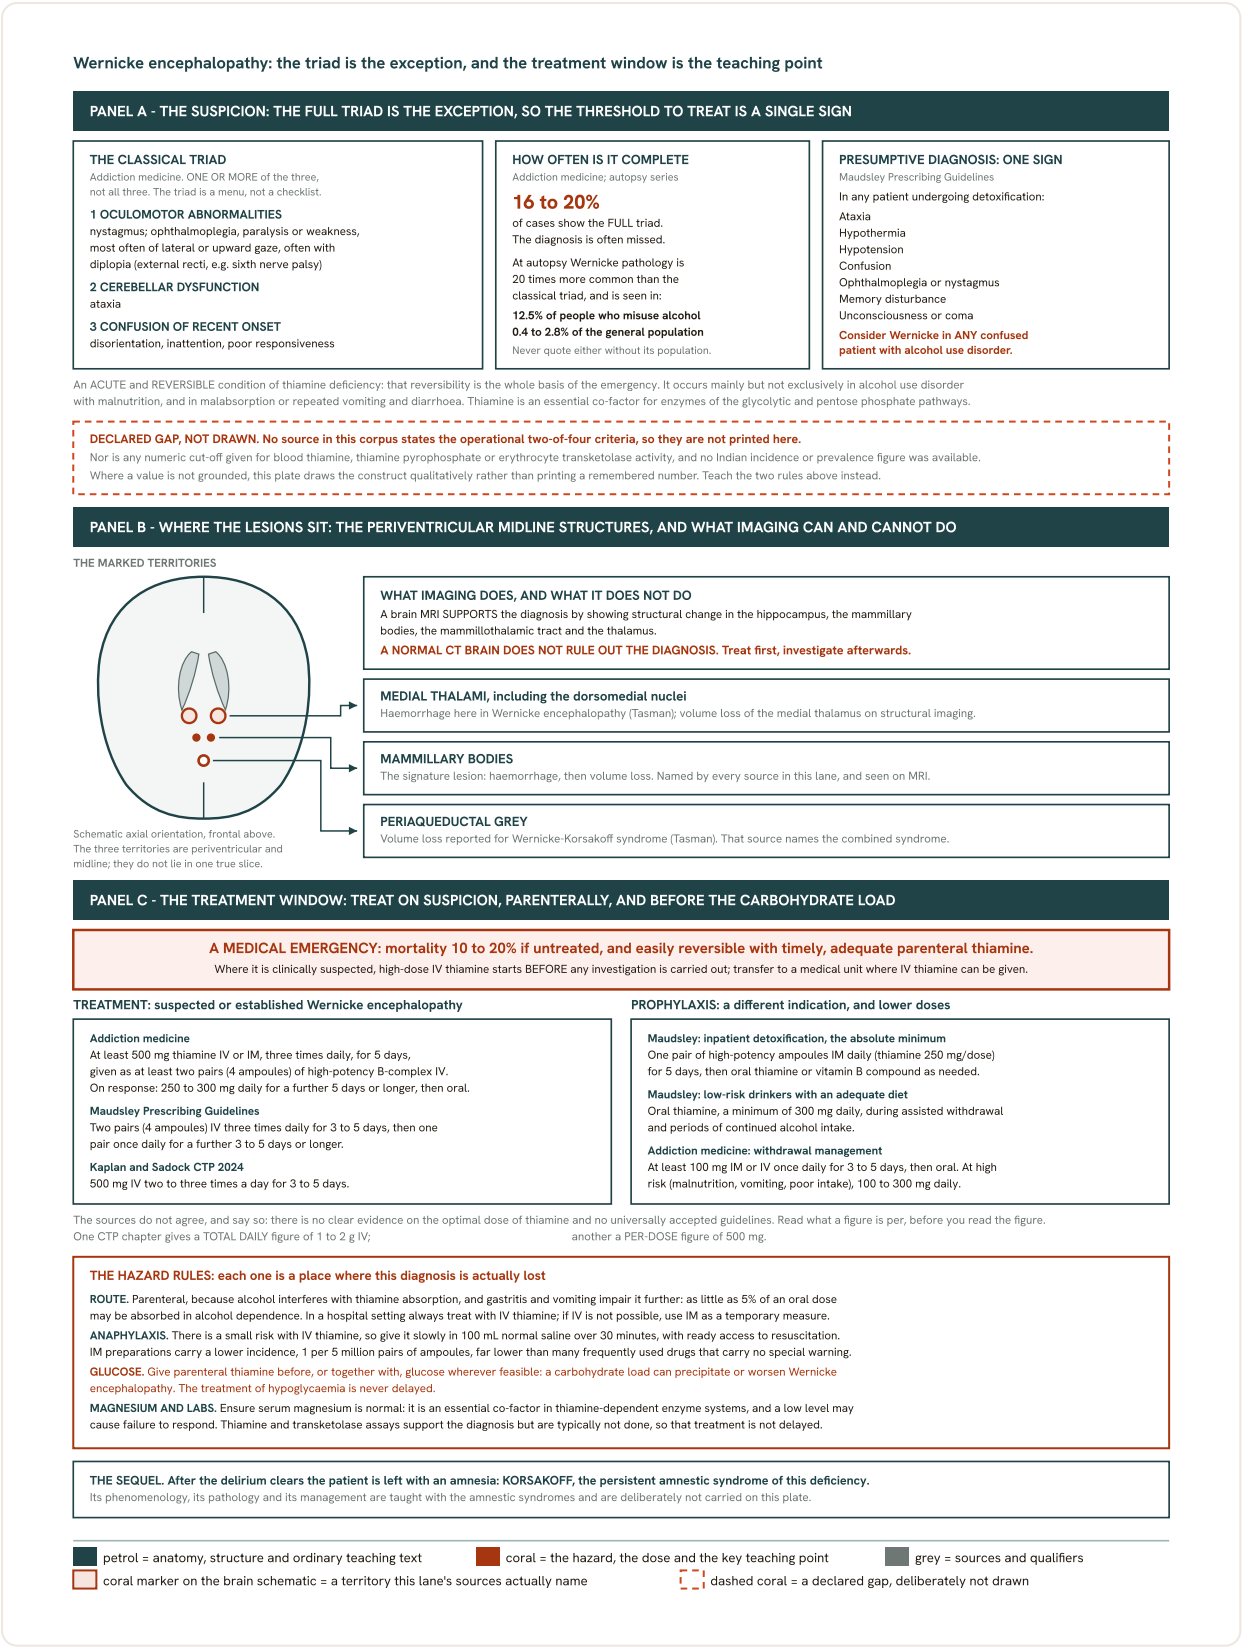


Fig 33.5 Wernicke encephalopathy: the classical triad, the periventricular midline territories, and the thiamine treatment window. The complete triad is the exception, so the threshold to treat is a single sign in a patient at risk, and treatment precedes investigation. Doses differ between sources and the sources say so; read the unit and the indication before the number, and never carry a prophylaxis dose across to treatment. Korsakoff, the persistent amnestic sequel, is named here and taught with the amnestic syndromes. (Addiction Medicine, Saunders et al.; The Maudsley Prescribing Guidelines; Kaplan and Sadock CTP 2024; Tasman, Psychiatry.)

12.4 Who to suspect it in, and the presumptive rule that replaces the triad

The operational threshold is therefore very low. Addiction medicine holds that there should be a high index of suspicion of Wernicke's encephalopathy and Wernicke-Korsakoff syndrome in patients with alcohol use disorders, and that the diagnosis should be considered in any confused patient with an alcohol use disorder until proved otherwise. The Maudsley Prescribing Guidelines convert suspicion into an explicit single-sign rule. Noting that the classical triad is rarely present and the syndrome is much more common than is recognised, they hold that a **presumptive diagnosis** should be made in any patient undergoing detoxification who experiences any one of the following:

- Ataxia, or ophthalmoplegia or nystagmus.
- Confusion, or memory disturbance.
- Hypothermia, or hypotension.
- Unconsciousness or coma.

Any history of malnutrition, recent weight loss, vomiting or diarrhoea, or peripheral neuropathy should also be noted. Hypothermia and hypotension are not in the triad at all, which is why the observations chart matters on a detoxification ward as much as the eye signs. The risk factors in alcohol dependence are acute withdrawal, malnourishment, decompensated liver disease, attendance at an emergency department, hospitalisation for a co-morbid condition, and homelessness. Three of those six are contextual rather than clinical: the syndrome concentrates where care is fragmented.

Kaufman's Clinical Neurology for Psychiatrists notes that thiamine deficiency, whatever its cause, generally leads to absent deep tendon reflexes and loss of position sensation, deficits that may remain asymptomatic unless the patient walks in the dark; in the Wernicke-Korsakoff syndrome, amnesia, dementia and cerebellar degeneration accompany the neuropathy associated with alcoholism. Hutchison's Clinical Methods tabulates the same vitamin B1 signature: nystagmus, sixth cranial nerve palsy and ataxia, with symmetrical peripheral neuropathy as the peripheral counterpart. The dementia syndrome and its workup belong to the Dementias and neurocognitive disorders notes.

12.5 Investigations: what confirms it, and why you will not wait

Addiction medicine is explicit about the thiamine assays: low blood thiamine or thiamine pyrophosphate levels and low **erythrocyte transketolase activity** support the diagnosis, but these are typically not done, because it is important not to delay treatment while awaiting the results. No numeric cut-off or reference range for any of the three appears in the standard references drawn on here: know the tests by name and by direction of abnormality, and quote no threshold value.

One laboratory value does change management. Serum magnesium should be checked and normalised, because magnesium is an essential co-factor in many thiamine-dependent enzyme systems and low magnesium levels may result in failure to respond to thiamine therapy.

PITFALL

Escalating the thiamine dose in a non-responder without ever checking a magnesium level. The reasoning feels sound: right diagnosis, therefore too low a dose. But in a malnourished drinker with vomiting and diarrhoea, hypomagnesaemia is close to expected, and non-response is not evidence against the diagnosis.

12.6 Three orders that must not be reversed

The route is not negotiable, and the reason is pharmacokinetic. In a hospital setting always treat with intravenous thiamine, because alcohol interferes with thiamine absorption and the medical complications of drinking, such as gastritis and vomiting, further impair it. As little as **5%** of an oral dose may be absorbed in alcohol dependence. If the intravenous route is not possible, use intramuscular thiamine as a temporary measure. Kaplan and Sadock supply the pharmacodynamic half: high-dose intravenous thiamine is currently considered superior to oral supplementation in preventing and treating Wernicke's, because high serum concentrations are needed to drive thiamine across the blood-brain barrier.

The second order is the one every examination tests. It is imperative that parenteral thiamine is given before a dextrose drip, because a carbohydrate load has the potential to precipitate or exacerbate Wernicke's encephalopathy. The Maudsley makes it a standing rule: as thiamine is required to utilise glucose, a glucose load in a thiamine-deficient patient can precipitate the encephalopathy, and parenteral B-complex must be given before glucose in all patients presenting with altered mental status. That clause is not restricted to known drinkers: those who most need it are the ones whose history is not yet known. Read both statements for what they are, which is an ordering rule. What each source fixes is which of the two goes in first, not whether the other is given at all. The third order was given above: treatment precedes investigation.

Note the limit of what these sources address. Both state the ordering rule and stop there. Neither takes up the patient who is at once thiamine-deficient and hypoglycaemic, and no claim drawn on in this fragment says what to do when a documented low glucose and an undrawn thiamine make their demands at the same moment. That silence is a gap in the sources, not a permission granted by them: nothing quoted above licenses withholding or delaying the treatment of hypoglycaemia, and a candidate should not read the ordering rule as though it did. Answer the examination question with the order the sources give. The competing case is a matter for acute medical practice and is settled by no sentence printed here.

MNEMONIC

Parenteral before oral, thiamine before glucose, treatment before tests. Sequence, not omission: these sources are silent on the patient who is also hypoglycaemic, and nothing in them permits delaying treatment of a documented low glucose.

RULE **Suspicion, not confirmation, is the threshold for treatment.** The triad is complete in a minority, autopsy prevalence far exceeds clinical diagnosis, untreated mortality is substantial, and the treatment is a water-soluble vitamin. Both prescribing sources therefore set the threshold at a single sign in a patient at risk rather than at a criterion count. Treat first, parenterally, before glucose and before imaging. Before is a statement about sequence, not about omission: it fixes which of two things goes in first, not whether the other is given.

12.7 Treating it: locus, focus and modus, and three regimens that disagree

LOCUS. The diagnosis is made wherever the confused drinker presents: emergency department, acute medical ward, consultation-liaison referral, detoxification unit. Treatment has one location: if Wernicke's encephalopathy is suspected, the patient should be transferred to a medical unit where intravenous thiamine can be administered. A psychiatric ward that cannot give intravenous thiamine with resuscitation cover is not a treatment location.

FOCUS. First hours: the encephalopathy itself, confusion, eye signs and ataxia, aiming at complete reversal. Next days: consolidating that response and correcting the co-factor deficiencies that limit it. Longer term: preventing relapse into deficiency and the amnestic sequel.

MODUS. The three standard references agree on the magnitude of parenteral dosing and on loading followed by continuation, and disagree on frequency and duration.

SOURCE	REGIMEN FOR SUSPECTED OR ESTABLISHED WERNICKE ENCEPHALOPATHY	WHAT DIFFERS
Addiction medicine	At least 500 mg of thiamine intravenously or intramuscularly, three times daily, for 5 days, given in the United Kingdom as at least two pairs of ampoules, that is four ampoules, of high-potency B-complex vitamins.	Frequency fixed; duration flat.
The Maudsley Prescribing Guidelines	On an acute medical ward, at least two pairs of high-potency intravenous ampoules, that is four ampoules, three times daily for 3 to 5 days, then one pair once daily for a further 3 to 5 days or longer, until no further response is seen.	Frequency fixed; duration a range.
Kaplan and Sadock, substance-use chapter	500 mg of intravenous thiamine two to three times a day for 3 to 5 days, or 100 mg of oral thiamine a day usually given for many months.	Frequency given as a band; a per-dose figure in milligrams.
Kaplan and Sadock, amnestic-disorders chapter	High doses of thiamine, 1 to 2 g intravenously daily.	A total daily figure in grams.

The continuation phase is more consistent. If it responds, Addiction medicine continues with 250 to 300 mg of thiamine daily for another 5 days, or longer if improvement continues, followed

by oral thiamine and multivitamin supplementation thereafter and as an outpatient, adding that for patients with enduring ataxia, polyneuritis or memory disturbance, high-potency vitamins should be given for as long as improvement continues. That rider makes continued improvement, not the calendar, the stopping rule.

TRAP **Answering as though one of these regimens is the regimen.** A stem offering four dose options and expecting one is testing a consensus that does not exist. Addiction medicine says there is no clear evidence on the optimal dose and no universally accepted guidelines; the Maudsley notes insufficient randomised evidence as to best dose, frequency or duration. One textbook even holds two formulations, milligrams per dose in one chapter and grams per day in another. Do not average them or silently pick one: state the magnitude, give each source's frequency and duration, and name the absence of consensus.

12.8 Prophylaxis, and a minimum that two sources set differently

Most thiamine given in psychiatry is preventive, and dosing is tiered by risk and route.

INDICATION	REGIMEN AS STATED	SOURCE
Prophylaxis, low-risk drinker with adequate diet and no neuropsychiatric complications	A minimum of 300 mg of oral thiamine daily during assisted alcohol withdrawal and periods of continued alcohol intake.	The Maudsley Prescribing Guidelines
Prophylaxis, inpatient alcohol detoxification	One pair of high-potency intramuscular ampoules daily, containing thiamine 250 mg per dose, for 5 days, then oral thiamine or vitamin B compound as needed; stated as the absolute minimum for all inpatients.	The Maudsley Prescribing Guidelines
Prophylaxis, alcohol withdrawal management	At least 100 mg of thiamine intramuscularly or intravenously once daily for 3 to 5 days then orally, or at least 300 mg of thiamine per day orally, maintained on oral thiamine 100 mg daily with multivitamins.	Addiction medicine
Prophylaxis, patient at high risk of Wernicke encephalopathy	In malnutrition, vomiting or poor oral intake, 100 to 300 mg daily of intramuscular or intravenous thiamine, with the warning that these are prevention doses and higher doses are used for treatment.	Addiction medicine

Both sources cost a pair of high-potency ampoules at the same 250 mg of thiamine per dose, so the disagreement is about the recommended floor, not the product. Their oral arms agree exactly.

TRAP **Quoting a prophylactic minimum as though there were one, or using a prevention dose to treat.** The Maudsley calls its inpatient parenteral regimen an absolute minimum for all inpatients; Addiction medicine calls its own, materially lower per dose, a minimum too. Both are grounded, so a stem asking for the minimum parenteral prophylactic dose has more than one defensible answer. The second error is worse clinically: a patient with a single presumptive sign has crossed into the treatment column, and the dose must cross with them.

IN INDIA

No Indian incidence or prevalence figure for Wernicke's encephalopathy appears in these standard references, so every epidemiological number above comes from Western autopsy and clinical series and should not be silently naturalised. Both prescribing sources dose in pairs of a United Kingdom high-potency parenteral B-complex preparation and name no Indian equivalent. Prescribe in milligrams of thiamine with route and frequency written out, using the sources' own figure of one pair containing thiamine 250 mg per dose, rather than an ampoule count the ward stock will not match. A normal computed tomography scan does not exclude the diagnosis, which weighs more where magnetic resonance imaging means a wait, a transfer and an out-of-pocket cost. Give the thiamine while the scan is being arranged.

12.9 The boundary: the reversible window and its irreversible sequel

Everything above earns its urgency from one sentence about time. Addiction medicine states that most patients who do not recover promptly within the first **48 to 72 hours** develop Korsakoff's syndrome with irreversible memory loss. Note what the source does not say: whether the window runs from symptom onset or from the start of thiamine. Neither should a candidate. It fixes the scale, days rather than weeks.

On the far side sits **Korsakoff's syndrome**, described as a persistent alcohol-induced amnestic confabulatory neurocognitive disorder and a sequel to untreated or inadequately treated Wernicke's encephalopathy, with an estimated 20% of affected individuals requiring long-term institutionalised care. Undertreatment is what the source blames: an oral prescription or a course stopped early is itself a mechanism of the chronic syndrome. The Maudsley states it bluntly: if untreated, Wernicke's encephalopathy progresses to Korsakoff's syndrome. Tasman adds the nosological caution that although the term is often used for all cases of chronic anterograde amnesia with a retrograde component, it is properly reserved for the form occurring secondary to thiamine deficiency. The syndrome in full belongs to the Stroke, TBI, delirium and focal brain syndromes notes.

The boundary is pharmacological as well as temporal: high-dose intravenous thiamine is likely to partially or completely reverse Wernicke encephalopathy, but its benefits in the chronic stages are not established. That asymmetry is why the acute intervention is framed as prophylaxis against the irreversible sequel: patients in alcohol withdrawal with suspected Wernicke encephalopathy should be given high-dose thiamine intramuscularly or intravenously, since oral absorption is very poor, to prevent progression to Korsakoff syndrome, and more broadly prevention of cognitive impairment is possible in those at risk of thiamine deficiency, from long-standing

alcohol misuse or severe malnutrition, through vitamin B1 given particularly in intravenous form.

Symptoms of Wernicke's encephalopathy may be mixed with those of Korsakoff's syndrome as a result of repeated episodes of clinical or subclinical thiamine deficiency, which is why the combined term **Wernicke Korsakoff syndrome** exists at all. The boundary is therefore a gradient rather than a line, and each unremarked episode moves the patient along it before anyone has recorded a sign.

TRAP **Quoting a single recovery figure for the Korsakoff amnestic syndrome.** Kaplan and Sadock's amnestic-disorders chapter states that about 25% of patients with Korsakoff syndrome recover, while the rest continue to suffer varying degrees of memory impairment and a significant minority need permanent placement, and Tasman corroborates it exactly: on long-term follow-up about one-quarter eventually recover, half show some improvement, and in the remaining quarter none is seen. The same textbook's substance-use chapter, however, calls the syndrome permanent in at least a partial form in about 50% of the people affected, implying a very different recovery rate. Two of the three readings agree, so print that distribution and name the substance-use chapter's wording as the discrepant reading rather than suppressing it.

GOOD TO REMEMBER

Do not let the anaphylaxis warning become a reason to withhold or to delay. There is a small risk of anaphylaxis with intravenous thiamine, so slow administration in 100 mL of normal saline over 30 minutes is recommended, with ready access to resuscitation measures; intramuscular preparations carry a lower incidence than intravenous ones, at 1 per 5 million pairs of ampoules, far lower than many frequently used drugs that carry no special warning.

Wernicke's encephalopathy is an acute, reversible thiamine-deficiency emergency in which the complete triad appears in a minority, so treat any confused patient at risk on a single sign with high-dose parenteral thiamine, before glucose and before imaging, check the magnesium if there is no response, and note that most of those not recovering promptly within the first 48 to 72 hours go on to the permanent amnestic syndrome thiamine can no longer reverse. Before means sequence, not omission: on the patient who is at once thiamine-deficient and hypoglycaemic the sources are silent, and nothing here licenses delaying treatment of a documented low glucose.

13 · Vitamin B12 and folate deficiency: cognitive and psychiatric manifestations

These two vitamins are the reason a psychiatrist orders blood tests at all. They are the standing example of a treatable cause of psychiatric illness hiding inside an ordinary presentation. Both are cofactors in the same stretch of one-carbon chemistry, both produce a macrocytic marrow, both produce cognitive and affective change, and only one of them damages the spinal cord. Everything examinable descends from that asymmetry.

The marks sit in four places: a phenotype wide enough to fit almost any primary psychiatric disorder; a laboratory of non-interchangeable thresholds and a confirmatory metabolite; the

discrimination between the two deficiencies; and treatment, where the classic error is a right drug given alone.

200-250 pg/mL	400 pg/mL	1.5% TO 15%	ABOUT 20%
B12 LABORATORY LOWER LIMIT	SCREEN-FURTHER THRESHOLD	B12 DEFICIENCY, OLDER ADULTS	FOLATE INSUFFICIENCY, US SURVEY DATA

13.1 Two vitamins, one stretch of chemistry

Folate matters to psychiatry entirely as an enzymatic cofactor, and the processes it serves are neuropsychiatric ones: neurotransmitter synthesis, central nervous system phospholipid synthesis, and homocysteine metabolism. B12 works alongside it in the same reactions, which is why the two deficiencies overlap so heavily in the blood and in the mind. That overlap is the mechanism, not a nuisance.

When B12 is absent the reactions it serves stall and their substrates accumulate, so that both homocysteine and methylmalonic acid rise. That is the entire basis of confirmatory testing: a metabolite piling up behind a blocked step is a more faithful witness to functional deficiency than the circulating vitamin. And because folate can substitute for B12 in one reaction but not the other, giving folate alone repairs part of the picture and lets the rest progress.

The supply routes differ, and the difference is load-bearing. B12 is essentially found only in animal products, and deficiency can arise either from decreased intake or from failure of absorption, since intrinsic factor is a gastric protein required to facilitate dietary absorption. One dietary class plus an obligatory gastric chaperone gives two independent failure modes, and a history establishing one does not exclude the other. Folate has no such chaperone, so its deficiency is driven by intake, demand and drugs rather than anatomy.

13.2 Who becomes deficient

Learn the two risk profiles separately, because they overlap less than the clinical pictures do. For B12, risk rises with older age, a vegan or vegetarian diet, pernicious anaemia, the malabsorption syndromes, and chronic use of metformin, proton-pump inhibitors and histamine-2 receptor antagonists. The drug triad deserves attention: all three are prescribed long term, for reasons unrelated to the nervous system, and none is likely to be flagged as a neuropsychiatric risk. The malabsorptive causes group anatomically.

- Gastric bypass surgery
- Inflammatory bowel disease
- Celiac disease
- Small intestinal bacterial overgrowth
- Ileal resection

The yield of testing older patients is real: estimates of B12 deficiency in older adults range from 1.5% to 15%, and low B12 status has been linked to depression and cognitive dysfunction. Treating this as a disease of the elderly is now questionable, because deficiency may be increasing in younger adults, and the Framingham Offspring Study found similar deficiency rates across the

age bands 26 to 49 years, 50 to 64 years, and 65 years and older. Read that precisely: it does not extend the prevalence band to the young, it says the assumed age gradient was absent in that dataset.

Folate runs on different drivers: inadequate diet, excessive alcohol intake, women of childbearing age, non-Hispanic black women, pregnancy, dialysis, celiac disease, inflammatory bowel disease, and methylenetetrahydrofolate reductase gene mutation. Two psychiatric populations are picked out there without being named: low folate levels are found in patients who abuse alcohol or who have poor nutrition, which covers much of any general caseload, and two drugs are specified: methotrexate and phenytoin can both lead to folate deficiency. The methotrexate link is firm enough that a standard clinical-methods text lists it as the named deficiency syndrome for folic acid, with macrocytic anaemia, thrombocytopenia and a megaloblastic bone marrow as the principal signs.

13.3 The psychiatric presentation

The psychiatric syndrome of B12 deficiency has no signature. What it has is breadth: depression, irritability, insomnia, cognitive slowing and psychosis. Every item is an ordinary reason for referral and no combination of them points at a vitamin. Folate produces a narrower version, with cognitive impairment and depression as its symptoms, and low folate levels associated with major depressive disorder, dementia and schizophrenia. Across both vitamins the nutritional deficiencies have been implicated in dementia, delirium, psychosis, fatigue and personality changes.

A list that wide is useless as a prompt and valuable as a rule. If no phenotype triggers the test, something else must: age, the drug list, the diet, or policy for unexplained presentations. Delirium as a syndrome belongs to the Stroke, TBI, delirium and focal brain syndromes notes; what this topic supplies is the correctable nutritional cause sitting behind a share of it.

> **GOOD TO REMEMBER**
>
> Personality change is the item most often handled as a formulation rather than a symptom. Relatives describing a patient as irritable, apathetic and unlike themselves over months, with disturbed sleep and slowed thinking, present a shape that attracts an interpretive explanation more readily than a blood test. Either vitamin can produce it.

13.4 The nervous system and the blood: what B12 damages

The neurology is what separates B12 from folate, and it is best held anatomically. B12 deficiency affects multiple levels of the nervous system, and may present as peripheral neuropathy, myelopathy, or a cognitive syndrome. The underlying lesion is demyelinating: B12 deficiency causes demyelination in the posterior tracts of the spinal cord, in the cerebrum and in other areas of the central nervous system, and it causes dementia and a peripheral neuropathy. Name the tract and you have the cord syndrome, the neuropathy, and the reason the two coexist.

The named cord syndrome is **subacute combined degeneration of the spine**, which travels in the standard tables with **macrocytic anaemia** under the deficiency syndrome **pernicious**

anaemia. That triad is the classical teaching picture, correct as far as it goes; its limitation is the subject of the next section.

Absorptive failure dominates in practice. B12 deficiency is usually the result of poor absorption due to pernicious anaemia or gastric disease, the clinical expression of the point about **intrinsic factor**. It also predicts something counter-intuitive: if the lesion is absorptive, oral repletion ought to fail. Why it does not is taken up under management.

13.5 Cognitive decline, and how far it reverses

This is where the received teaching is least reliable. Two propositions are carried forward together and both need qualifying. The first is that macrocytosis is the tripwire. It is not: B12 deficiency is often suggested by the presence of a macrocytic anaemia, but B12 deficiency dementia may occur in the absence of both macrocytosis and anaemia. The same holds on the other side, where folic acid deficiency is a very rare cause of dementia which likewise can occur in the absence of anaemia. A normal full blood count excludes neither, and this is the commonest route by which the diagnosis is lost.

The second proposition is reversibility, more heavily qualified than most candidates expect. Although neurologists often cite B12 deficiency as the quintessential cause of "reversible dementia", several studies have found that the dementia resists B12 repletion. Hold both halves: the disease belongs on every list of treatable causes of cognitive decline, but the promise that treating it restores cognition is not one the literature fully honours. That argues for testing early, not against testing. The dementia syndrome and its workup belong to the Dementias and neurocognitive disorders notes, and cognitive instruments to the Psychotherapies, clinical assessment, MSE and nosology notes; this topic owns the disease, its cognitive signature and the limits of its reversibility.

> **PITFALL**
>
> Ordering a full blood count as the screen for B12 deficiency in a patient with cognitive decline. The error is logical: the classical association is haematological, and a normal mean corpuscular volume feels like reassurance. It is not, because the deficiency dementia occurs without macrocytosis and without anaemia. If the question is neuropsychiatric, the assay has to be the vitamin.

13.6 Reading the laboratory: which threshold does which job

Marks are lost here by treating several different numbers as one cut-off. B12 status is measured in serum or plasma, with **200 to 250 pg/mL** as the lower limits used by most laboratories, while some suggest that levels up to 400 pg/mL should be screened further for deficiency. The higher figure exists precisely because a value inside the reference interval does not exclude functional deficiency.

Folate has its own set. Deficiency is defined as a serum folate below **3 ng/mL**, while **folate insufficiency** sits at approximately below 7 ng/mL, a level that matters because insufficiency in women of childbearing age is a risk factor for neural tube defects in offspring. A pharmacology reference adds the exclusion figure and the red cell measure.

VALUE	WHAT IT IS	WHAT IT IS NOT
B12 200 to 250 pg/mL	Laboratory lower limit of normal	Not a functional deficiency threshold
B12 400 pg/mL	Threshold for screening further	Not a diagnosis, not a treatment trigger
Serum folate below 3 ng/mL	The definition of folate deficiency	Not applicable to red cell folate
Serum folate approximately below 7 ng/mL	Folate insufficiency, carrying the neural tube defect risk	Not the deficiency definition
Serum folate above 5.0 ng/mL	Level above which folate deficiency can be excluded, though 2% to 5% of healthy adults fall below it	Not a target, not a deficiency cut-off
Red cell folate above 140 ng/mL	Reference range for chronic folate status, less affected by acute ingestion than the serum level but more time consuming and costly to measure	Not a deficiency cut-off; none is stated here

The last row carries an examinable absence: the serum figures cannot be transposed onto red cell folate, because the two measure different things over different timescales.

13.7 The metabolic intermediaries, and what each discriminates

When the vitamin level is equivocal the question moves from concentration to function. Sensitivity in determining B12 deficiency may be enhanced by measuring intermediate metabolites such as methylmalonic acid and serum homocysteine. The indication is specific: they are useful particularly when vitamin levels are in a borderline range where deficiency cannot definitely be ruled in or out. They are not a first-line screen.

The discriminating rule is the one thing here that must be automatic. B12 and red blood cell folate levels constitute the screening tests, followed up by homocysteine and methylmalonic acid: homocysteine is elevated in both B12 and folic acid deficiency, while methylmalonic acid is elevated in B12 deficiency alone. A pharmacology reference states the same independently, that **serum homocysteine** is elevated in both while **serum methylmalonic acid** is elevated only in B12 deficiency, and it follows mechanically from the coenzyme chemistry above.

MNEMONIC

Both rise together; methylmalonic rises alone

Homocysteine is the shared metabolite and rises in either deficiency, so it says something in the pathway is missing but not which. Methylmalonic acid sits behind a step only B12 serves, so a rise points at B12 specifically.

For the confirmatory figure, one source gives an explicit boundary: in borderline cases methylmalonic acid is recommended as a more sensitive measure of B12 status, with an elevated

level above **0.271 micromol/L** serving as a reliable cut-off indicating deficiency. A second standard reference prints an upper reference limit for the same analyte falling immediately below that value.

TRAP **Reading a reference limit as a diagnostic cut-off because the two numbers coincide.** Kaplan and Sadock give a cut-off for methylmalonic acid, above which deficiency is indicated. Goodman and Gilman give an upper reference limit for the same analyte, landing in essentially the same place. The coincidence looks like corroboration. It corroborates a number, not a category: a reference limit comes from a healthy population and answers whether a result is unusual, a cut-off is derived against a diagnosis and answers whether it means deficiency. A stem offering one and asking for the other is testing exactly this. A second layer: the printed units for these metabolite ranges differ between the two texts by orders of magnitude, a typesetting artefact rather than a disagreement about biology.

No numeric homocysteine threshold appears in this chapter, and the omission is deliberate: the available reference figures carry units irreconcilable with any clinical assay as printed, and the alternative candidate value sits in a text fragment whose subject cannot be established. Learn the direction, not a number.

13.8 Management: repletion, the masking trap, and folate in depression

LOCUS. Almost all of this is outpatient: detection in clinic while investigating cognitive change, treatment-resistant depression or an unexplained psychosis, and repletion in clinic too. Inpatient care is driven by the psychiatric syndrome rather than the vitamin, principally the psychotic and confusional presentations. There is no emergency limb here as the encephalopathies have one, but there is a time limb: the cord and the cognitive syndrome do not wait for a convenient appointment.

FOCUS. Acutely the target is the deficiency, not the psychiatric symptom, which is downstream. In the short term it is the cause, since a dietary deficiency, an absorptive failure and a drug effect need different follow-ups. In the medium term it is the neurological deficit, watched with the reversibility caveat in mind. In the long term it is maintenance: an absorptive lesion does not resolve, and a patient who stays on the offending drug stays at risk.

MODUS. Two evidence anchors carry the pharmacological limb, both frequently misquoted. For repletion in vitamin B12 deficiency, a 2018 Cochrane review graded the evidence low quality and found that daily oral supplementation at **1,000 to 2,000 mcg per day** is as efficacious and safe as, and potentially more cost-effective than, intramuscular supplementation. The grading is part of the finding. The parenteral route keeps a role where repletion must be rapid: rapid repletion of vitamin B12 is usually accomplished by intramuscular administration daily for a week, then weekly for the first month, followed by monthly maintenance therapy. The schedule is the examinable part; the quantity attached to it in the source cannot be reconciled with any plausible clinical dose and is therefore not printed rather than silently corrected from memory. The paradox of treating an absorptive disease by mouth resolves cleanly, since chronic oral repletion is possible even in patients with low intrinsic factor activity, through a low-affinity uptake mechanism.

RULE **Never replete folate in a patient whose B12 status is unknown.** In the absence of B12, which acts as a coenzyme, homocysteine and methylmalonic acid accumulate. Administration of folate, which also acts as a coenzyme, alleviates some features of B12 deficiency such as the megaloblastic anaemia; but without concomitant administration of B12, folate will not prevent subacute combined degeneration. The clinical shape is what makes it dangerous: the blood count improves, the patient looks treated, and the cord goes on demyelinating underneath a corrected haemoglobin. Treat a haematological response as no evidence at all about the nervous system.

Folate has a second, distinct indication, and the two must not be run together. As augmentation in major depressive disorder, a meta-analysis of nine randomised controlled trials with 1,061 participants found folate adjunctive to antidepressant therapy improved depressive symptoms in adults with major depressive disorder with a small effect size, and raised the remission rate, with remission induced in 48% of the folate group against 27% on placebo. The strategies used were folic acid 0.5 to 10 mg per day and **L-methylfolate** at 15 mg per day, over a mean study duration of 15.2 weeks. Antidepressant pharmacology and the sequencing of augmentation belong to the Mood disorders: pharmacotherapy notes; what belongs here is that this is a nutritional agent used pharmacologically and it is not risk-free, since L-methylfolate at 15 mg per day for depressive augmentation has triggered irritability and agitation in a few reported cases, one of them in a patient with bipolar disorder type 1.

There is an upper bound, and its rationale closes the loop with the masking trap: the Food and Nutrition Board at the National Academies of Sciences, Engineering, and Medicine has set an upper limit for supplemental folate in adults over 19 years, one stated reason being that folate can mask B12 deficiency. The quantity is not printed here. The figure as the source gives it carries the same order-of-magnitude unit problem already flagged for the parenteral B12 dose, and a gram-for-milligram slip overstates a safety ceiling a thousandfold, so this limit is carried as a principle and its value read off a current dietary reference intake table rather than reproduced on trust. What is examinable is the existence of a ceiling and the reason for it, which is the masking rule again in another guise. The psychological and social modes are brief but not trivial, since adherence to lifelong maintenance is a different problem from adherence to a short course.

IN INDIA

Two numbers on this page are population-bound. The rarity of frank folate deficiency in the source literature is a post-fortification finding: folate deficiency is rare in the United States because of the 1996 Food and Drug Administration mandate requiring folate supplementation of enriched grain products, completed in 1998, and the insufficiency rate quoted earlier comes from that same fortified population. Neither transfers to a population without an equivalent long-standing nationwide grain-fortification programme, and quoting either as a general fact quotes a policy outcome as a biological one. The B12 side runs the other way: a vegan or vegetarian diet is a listed risk factor and B12 is essentially found only in animal products, so where lifelong vegetarianism is common, dietary deficiency is a far more available explanation than the absorptive causes dominating Western teaching. Practically, serum B12 is cheap and widely available across Indian laboratories while methylmalonic acid is a send-out in most centres at a cost the patient meets, so the borderline-value pathway is often settled on clinical grounds and a therapeutic trial instead. Where the metabolite is unaffordable, record that rather than filing a borderline B12 as normal.

Homocysteine rises in both B12 and folate deficiency while methylmalonic acid rises in B12 deficiency alone; B12 deficiency dementia occurs without macrocytosis and without anaemia, so a normal blood count excludes nothing; and folate given without B12 corrects the anaemia while the cord degenerates.

14 · Pellagra and other vitamin deficiency syndromes

These are the diagnoses where being right still changes the outcome. A patient whose cognition, mood or conduct is deteriorating because a vitamin is missing recovers when it is replaced and not when it is not, and the window is finite. Two features make the group harder than it looks: each syndrome is named after its complete, textbook form, which is the form you will least often meet, and deficiencies travel in company, so the vitamin you identify is seldom the only one absent.

6D85.8	15 DAYS	UNDER 20 NG/ML	SMD 0.58
ICD-11 DEMENTIA DUE TO PELLAGRA	HALF-LIFE OF 25(OH)D	VITAMIN D DEFICIENCY	VITAMIN D IN UNIPOLAR DEPRESSION

14.1 Pellagra: the name, the triad and the third D that keeps moving

Pellagra takes its name from the skin, not the brain: the word is Italian for rough skin, from pelle for skin and agro for rough. The standard neurology text for psychiatrists calls it a starvation-induced disorder, which locates it among conditions of deprivation, not primary neurological disease.

MNEMONIC

Dementia, dermatitis, diarrhoea

The three Ds. An aide-memoire for which three systems to examine, not a checklist to satisfy before the diagnosis can be entertained. Skin, gut and brain are involved because niacin sits in the energy metabolism of tissues with high turnover.

The third D is the one to watch, because the sources do not print the same word. The neurology text gives dementia in its own account of the triad and, in a worked clinical answer elsewhere in the same book, dermatitis, diarrhoea and delirium; the psychiatric glossary sidesteps the choice and defines the syndrome as diarrhoea, dermatitis and mental disturbance. These are not three competing facts but one fact at three levels of resolution. That book's own neuropsychiatric shorthand puts both words side by side, giving pellagra with delirium and dementia, because niacin deficiency produces a diffuse encephalopathy that can present acutely and fluctuantly, or insidiously and persistently. Delirium as a syndrome and its workup belong to the Stroke, TBI, delirium and focal brain syndromes notes, the dementia syndrome and its investigation to the Dementias and neurocognitive disorders notes. An examiner writing dementia and one writing delirium are both quoting a real source.

14.2 The incomplete triad, and the skin that gives it away

Chronic niacin deficiency produces the classic picture through the gradual onset of the three Ds, with dermatitis on sun-exposed areas, but in most cases of **pellagrinous dementia** only one of the other Ds is present, as with the classical triad of thiamine deficiency: a general property of nutritional encephalopathies, not a quirk of one disease.

In most patients with pellagrinous dementia the triad is incomplete: only one of the other two Ds accompanies the cognitive syndrome, so a normal gut or normal skin excludes nothing.

RULE **These syndromes are diagnosed from context, not from completeness.** The diet, the alcohol history, the drug list and the socioeconomic circumstances carry more weight than the number of classical features present. Anyone who waits for the full triad before naming pellagra will miss most of the pellagra.

The skin is the best single clue when present. The rash is necklace-like: dermatitis of sun-exposed areas including **Casal's necklace** heads the clinical methods list of signs, which also includes dementia, poor appetite, difficulty sleeping and diarrhoea. Poor appetite and disturbed sleep matter for psychiatry: they are the complaints that route a malnourished patient into a psychiatric clinic with a provisional label of depression. The photosensitive distribution is the discriminator: a rash respecting the collar line, the backs of the hands and the face is describing sunlight, and sunlight plus cognitive decline plus a deprived diet is specific enough to act on.

14.3 Where it comes from: maize, tryptophan, alcohol and isoniazid

The historical account is dietary. Pellagra was endemic in the rural southeastern United States in the early 1900s and results from the absence of niacin from the diet; the population affected there subsisted on maize, which lacks niacin and also lacks its precursor, tryptophan. The second clause is the half candidates drop. A diet deficient in niacin can still hold if it supplies enough of the amino acid precursor, so a staple poor in both is far more dangerous than one poor in only one. **Tryptophan** is why niacin requirement is not a fixed quantity.

Three routes bring a modern patient to the same place.

- Deprivation and alcohol. The condition is described as typically found only in malnourished people with alcohol use disorder, and the mechanism is doubled: alongside dietary deficiency, niacin absorption is itself impaired in chronic alcohol dependence.
- Drugs. Isoniazid can also cause pellagra.
- Company. Other B vitamin deficiencies are often associated, and anaemia and hypoproteinaemia are common in pellagra; with nutritional deprivation, Wernicke-Korsakoff disease may be comorbid with it. Both of those have their own homes, in this chapter and in the Stroke, TBI, delirium and focal brain syndromes notes.

IN INDIA

The drug cause is the divergence that matters. Where anti-tuberculosis treatment is commonplace, isoniazid is not a curiosity at the end of an aetiology list but the exposure most likely to be in your patient's current prescription. The same drug is named in the clinical methods literature as the agent associated with pyridoxine deficiency, so one medication reaches into both halves of this section. Take a tuberculosis treatment history in any patient presenting with an encephalopathy, a photosensitive rash or a peripheral neuropathy. The second point is an absence: the sources used here give no Indian prevalence figure for pellagra or vitamin D deficiency, so no local rate is quoted.

14.4 The niacin attribution, and the argument that has not been settled

Almost every source treats pellagra and niacin deficiency as the same statement. The psychiatric glossary defines the disease as a deficiency of niacin, that is nicotinamide; the clinical methods table lists pellagra as the deficiency syndrome of vitamin B3, niacin, nicotinic acid; the general psychiatry text has chronic niacin deficiency causing pellagra. The international classification goes furthest, giving dementia due to pellagra its own code, the only nutritional deficiency so distinguished. When a diagnostic system carves out a code, it has committed to the aetiology.

One text dissents. The neurology reference for psychiatrists records that, despite pellagra's status as a classic illness, the role of niacin deficiency has been challenged, with deficiencies of other nutrients either co-existing with it or being more likely to be the actual cause. The same book hedges in its own definition, writing that niacin deficiency is associated with or causes pellagra, then elsewhere states flatly that the disease results from absence of niacin in the diet. The challenge is not even internally consistent within the book that raises it.

TRAP **Treating the niacin attribution as either settled beyond question or as discredited.** Both readings are available in the standard texts and both are wrong as absolute statements. The position that survives an examiner from either camp is that niacin deficiency is the established and codified account, that a serious source questions its sufficiency because these patients are rarely deficient in one nutrient alone, and that this is precisely why treatment is not niacin alone.

14.5 Managing pellagra

LOCUS. Most pellagra is treated where the patient already is: the general medical ward for the acutely encephalopathic or dehydrated, the psychiatric inpatient unit for those admitted for the mental state and found malnourished, the outpatient clinic for prophylaxis in someone at risk. What does not work is a psychiatric admission where nutrition is somebody else's problem.

FOCUS. Acutely, the target is the encephalopathy and the systemic deficit, alongside any comorbid thiamine-responsive illness. In the short and medium term it is the rash, the diarrhoea and the reversal of cognitive impairment, with the anaemia and hypoproteinaemia that accompany the disease. Long term it is the cause: an alcohol use disorder, a monotonous cereal-based diet, or a drug exposure that will reproduce the illness once replacement stops.

MODUS. The pharmacological mode is vitamin replacement and it is deliberately plural: the treatment of pellagra is niacin, and physicians usually add thiamine and other nutrients to the medical regimen: a patient malnourished enough to develop pellagra needs more than one vitamin. Preparations exist as **nicotinamide** as well as nicotinic acid, and the daily nutritional requirement of niacin, which is not a treatment figure, is given as **20 mg per day**. Without a change in diet, drinking or supervision, replacement is temporary.

On the treatment dose this chapter prints one source's figures and says so. For established pellagra the pharmacology text gives **200 to 500 mg per day** in divided doses, orally or parenterally, and for prophylaxis in people at risk 20 to 50 mg per day by mouth. No other reference consulted here states any niacin dose for pellagra: the neurology text says only that the treatment is niacin, and the general psychiatry and clinical methods texts give no figure. Treat those numbers as uncorroborated and check them against your formulary before prescribing.

PITFALL

Correcting one vitamin and considering the patient treated. The error is not a wrong diagnosis but an incomplete correction: niacin is given, the rash improves, the cognitive state improves partially, and the residual deficit is attributed to an established dementia or to alcohol-related brain damage. Since other B vitamin deficiencies frequently accompany pellagra and thiamine-responsive disease may be comorbid, a partial response is first a question about what else is missing and only later about irreversibility.

14.6 Pyridoxine: two clinical facts and one drug

This section is short because the evidence base is: the standard texts that spend pages on pellagra give pyridoxine two lines. Both are worth knowing exactly.

VITAMIN	DEFICIENCY SYNDROME	PRINCIPAL FEATURES NAMED
B3, niacin, nicotinic acid	Pellagra	Dermatitis of sun-exposed areas including Casal's necklace, dementia, poor appetite, difficulty sleeping, diarrhoea
B6, pyridoxine	Deficiency associated with isoniazid use	Poor appetite, lassitude, oxaluria, seborrhoea, neuropathy
D, ergocalciferol and cholecalciferol	Rickets and osteomalacia	Bone pain, proximal myopathy, waddling gait, growth retardation

The neuropsychiatric shorthand is the first fact: **pyridoxine deficiency** causes a neuropathy and seizures. Set beside the niacin sentence it gives a clean discrimination for a viva: niacin deficiency produces a diffuse encephalopathy with skin and gut involvement, pyridoxine deficiency a peripheral nerve and seizure threshold problem, without the rash and without a dementia.

The second fact is the drug association, and where the chapter deliberately stops. The clinical methods table names isoniazid as the cause and prints no prophylactic pyridoxine dose. The pharmacology literature consulted here carries a pyridoxine figure, but it is stated for sideroblastic abnormality rather than neuropathy prophylaxis, so it is the wrong indication as well as an unreliable read. No number is therefore given for pyridoxine, in either deficiency or prophylaxis: a plausible remembered dose in a revision text is how a wrong figure enters a resident's answer and stays there. Know the association, know that co-prescription with isoniazid is standard practice, and take the dose from your national programme guidance.

14.7 Vitamin D: the biomarker, the bands and the trap inside them

Vitamin D is the deficiency a psychiatrist is now most likely to be asked about, and it behaves differently from the B vitamins. Its named deficiency syndromes are skeletal, rickets in the growing skeleton and **osteomalacia** in the mature one, and the clinical methods table listing its features records no neuropsychiatric sign at all. The psychiatric interest comes from a separate literature about status and symptoms, not from the classical deficiency disease, and keeping those apart is most of the clarity available.

Status is measured as **serum 25-hydroxyvitamin D**, the optimal biomarker for two stated reasons: its relatively long half-life of 15 days, and the fact that it reflects endogenous production as well as vitamin D obtained from food and supplements. A short-lived marker would track this week's sunshine; a marker capturing only intake would miss the skin's contribution. The cut-points run below 20 ng/mL for deficiency, 20 to 29 ng/mL for insufficiency and above 30 ng/mL for sufficiency, with SI equivalents printed alongside.

TRAP **Quoting the SI band boundaries as though they were exact.** The conventional-unit bands tile the number line cleanly; their nanomole equivalents in the same source do not. Sufficiency is given as above 75 nmol/L while insufficiency is given as 50 to 70 nmol/L, so values between 70 and 75 nmol/L belong to no band, and the lower boundary is claimed by both insufficiency and deficiency depending on which line you read. This is one source being internally inconsistent rather than two bodies disagreeing: the ng/mL figures are primary and the nanomolar values are rounded conversions. Quote the conventional units as primary and never build an answer on a nanomolar boundary being exact.

Population data must stay in their own population. In a national survey of United States adults covering 2001 to 2010, one in three adults was considered vitamin D deficient and a further 40% met criteria for insufficiency. This chapter does not transplant those figures, which matters especially because dark pigmented skin is itself among the named risk factors. The rest of the list is directly usable in a psychiatric clinic: older age, obesity, physical inactivity, smoking, limited milk intake, limited sunlight exposure, fat malabsorption, primary hyperparathyroidism, chronic granulomatous disease, and medications including anticonvulsants, glucocorticoids and anti-HIV agents. Read it against a long-stay psychiatric population and it describes most of them.

14.8 Repletion, and what the psychiatric trials actually show

LOCUS for vitamin D is almost entirely outpatient, a clinic habit rather than a ward protocol. FOCUS moves in two phases, repletion to bring the measured level up and then indefinite maintenance, with the psychiatric symptom as a secondary target. MODUS is oral replacement, with either cholecalciferol, that is vitamin D3, or ergocalciferol, vitamin D2, being acceptable, alongside the social modes of sunlight exposure and diet. The guidance most commonly cited is the Endocrine Society's, whose 2011 recommendations are reproduced in the standard psychiatric textbook and set out below for established deficiency, not for routine supplementation of the replete.

GROUP BEING TREATED FOR VITAMIN D DEFICIENCY	REPLETION	DURATION	MAINTENANCE
Adults over 18 years	50,000 IU per week or 6,000 IU per day	Eight weeks in total	1,500 to 2,000 IU per day or 50,000 IU every 2 weeks
Children and adolescents aged 1 to 18 years	2,000 IU per day or 50,000 IU per week	At least six weeks in total	600 to 1,000 IU per day
Obese adults, malabsorption syndromes, or patients on drugs affecting vitamin D metabolism	6,000 to 10,000 IU per day	Titrated to the measured level	3,000 to 6,000 IU per day

Repletion is judged against a target rather than a symptom: the durations are estimated to achieve a level of at least 30 ng/mL, the same sufficiency threshold as before, which is why the biomarker is worth measuring twice.

GOOD TO REMEMBER

The third row of that table is a psychiatric row. The interfering medications named are anticonvulsants, glucocorticoids and other drugs that enhance activation of the steroid xenobiotic receptor, which results in destruction of 25-hydroxyvitamin D. A patient on long-term anticonvulsant treatment therefore needs the higher repletion and maintenance regimen and will drift back down on the standard adult dose. Anticonvulsants used as mood stabilisers are covered as drugs in the Mood disorders: pharmacotherapy notes; what belongs here is that their effect on vitamin D changes the dose you prescribe.

The psychiatric evidence is real but modest. In adults with unipolar depression, a meta-analysis of four randomised controlled trials with 948 participants found supplementation improved depressive symptom scores on the Hamilton Depression Rating Scale or the Beck Depression Inventory-II, with a moderate effect size of SMD 0.58, P less than .05. In attention deficit hyperactivity disorder, a separate meta-analysis of four randomised controlled trials with 256 children and adolescents aged 2 to 18 years found supplementation added to methylphenidate improved symptoms, with a small effect on total scores, SMD 0.39, and medium effects on inattention, SMD 0.67, hyperactivity, SMD 0.60, and behaviour, SMD 0.54, all P less than .01. Two meta-analyses, both of four trials, different participant numbers, different populations: keep each set of numbers attached to its own endpoint, because a merged figure is a fabricated one.

Enthusiasm is checked by pharmacology. Vitamin D is a **fat-soluble vitamin**, so excess supplementation carries a genuine risk of toxicity, and many participants in these trials were deficient or insufficient at baseline, so the benefit observed may be that of correcting a deficiency rather than of supplementing anybody. The stated responsible practice follows: check the serum 25-hydroxyvitamin D level in a patient with depression before starting supplementation, then follow evidence-based repletion and maintenance dosing.

14.9 When to test, and when to treat without testing

The last question is the one you will face on a ward round. The principle is stated plainly: it is the clinical history and examination that should provide the clues as to whether nutritional deficiency is contributing to the presentation. Testing follows suspicion; it does not generate it.

Where a deficiency is common and the consequence of missing it is severe, practice inverts. Thiamine and folate deficiency are common in patients who misuse alcohol, and although blood thiamine concentration can be measured, for patients at known risk most clinicians simply order thiamine supplementation without first ordering a laboratory evaluation. That is a rational asymmetry rather than sloppiness: the test is slow and the untreated disease is fast. Thiamine-responsive encephalopathy has its own section in this chapter, as do the cognitive and psychiatric consequences of vitamin B12 and folate deficiency; the persistent amnestic syndrome belongs to the Stroke, TBI, delirium and focal brain syndromes notes.

For the water-soluble vitamins in a deprived or alcohol-dependent patient, suspicion justifies treatment and the treatment is broad rather than targeted. For vitamin D, where the deficiency is chronic and the vitamin accumulates, you measure first and dose to the measurement. Reverse

that and you get the two characteristic errors here: a patient with an evolving nutritional encephalopathy waiting on an assay, and a patient with low mood taking a high-dose fat-soluble vitamin nobody ever measured.

The metabolic cost of treatment, diabetes and electrolytes

15 · Antipsychotic metabolic effects: weight gain, dyslipidaemia and glucose

This is where antipsychotic treatment collects its long bill. The test is not whether you know that antipsychotics cause weight gain, but whether you can separate three liabilities, attach each to the drugs and patients where it bites, and say where the evidence stops. **Both first-generation and second-generation antipsychotics** carry weight gain, dyslipidaemia, raised plasma glucose and diabetes, alongside hyperprolactinaemia, hip fracture, sexual dysfunction, extrapyramidal effects including neuroleptic malignant syndrome, anticholinergic effects, venous thromboembolism, sedation and postural hypotension. The class shorthand is that second-generation agents buy fewer extrapyramidal symptoms at the price of more metabolic harm: true as a class statement, useless as a prescribing one, since some older drugs behave like the safest of the new.

Two facts frame what follows. The burden is not evenly spread: on clozapine, nearly every exposed patient shows at least one of the triad of weight gain, insulin resistance and dyslipidaemia. And the baseline is not clean, because schizophrenia carries high rates of insulin resistance and diabetes and that was observed before antipsychotics existed. The drug amplifies an existing vulnerability rather than authoring it.

NEARLY 100%	42.7%	AS HIGH AS 61.6%	2 TO 3-FOLD
CLOZAPINE, AT LEAST ONE OF THE TRIAD	DIABETES, US CLOZAPINE DATA SET	METABOLIC SYNDROME, SAME POPULATION	DIABETES IN SCHIZOPHRENIA VERSUS POPULATION

15.1 Three liabilities, not one

The commonest error is to treat this as one process, with weight gain the engine and lipids and glucose the passengers. The cleanest counter-demonstration: olanzapine-treated and risperidone-treated patients gained similar weight, yet triglycerides rose four times as much on olanzapine, by 105 mg/dl against 32 mg/dl. Identical weight, very different lipids.

RULE **Weight, lipids and glucose are three separate liabilities and must be measured separately.** They cluster in the same high-risk agents, but antipsychotic-induced dyslipidaemia and antipsychotic-related insulin resistance can each occur without weight gain. A stable weight is reassurance about weight and nothing else.

15.2 Mechanism: impaired satiety, not a slowed metabolism

The mechanism of antipsychotic-induced weight gain, and of the two metabolic effects alongside it, is receptor-mediated and belongs here; antipsychotic action in psychosis and the principles of drug selection belong to the Schizophrenia: management, antipsychotics and adverse effects

notes, and the underlying receptor biology to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes. What matters here is how a receptor profile becomes a waistline.

The honest starting position is that **the mechanisms of antipsychotic-induced weight gain** are not well understood. The factors commonly implicated are 5-HT2C antagonism, H1 antagonism, D2 antagonism, and a rise in serum leptin producing leptin desensitisation. The most robust is **histamine H1 antagonism at hypothalamic sites regulating food intake**, with a possible added contribution from **5-HT2C antagonism**, supported by animal models and early polymorphism studies, though that genetic association is inconsistent.

The most useful mechanistic fact is negative. There is no evidence of any direct metabolic effect: weight gain results from increased food intake and, in some cases, reduced energy expenditure. The H1 account is one of impaired satiety, so patients eat more and metabolic rate is not switched off. That reframes the conversation, because satiety is something dietary structure and activity can engage with.

The mechanisms of antipsychotic-related diabetes are similarly unclear, and may include 5-HT2A and 5-HT2C antagonism, raised plasma lipids, weight gain itself and leptin resistance. Two agents differ in kind: olanzapine and clozapine appear to directly induce insulin resistance, which explains why impaired glucose tolerance appears on them in patients neither obese nor family-history positive.

15.3 Magnitude and time course of weight gain

Weight gain occurs primarily in the first few months but continues for many months or years afterwards, and in those who gain most the rate is maintained for at least two years. It is a long curve with a steep beginning, and the figures worth carrying are drug-bound and population-bound.

In a ten-year study of first-episode schizophrenia on mixed antipsychotics, mean weight gain was 15 kg, of which 9 kg came in the first year. That figure must never be transplanted onto a single drug or a chronic population. For clozapine, **60 to 75% of patients gain at least 4.5 kg during the first year**, and gain is often profound, with **4.5 kg in the first 10 weeks**. Naturalistic American data give a mean increase of 13.5 kg after 10 years on clozapine, of which that same early 4.5 kg came in the first 10 weeks. Two independent sources agree on it, which is why it is worth memorising.

Information on individual first-generation agents other than chlorpromazine and haloperidol is scarce, but one comprehensive analysis found first-generation drugs excluding haloperidol carried moderate risk, with 20 to 30% of people gaining more than 7% of original body weight in the medium term. Two things travel with that: more than 7% of original body weight is the conventional definition of clinically significant weight gain, and the population excludes haloperidol, which matters for the tier table.

GOOD TO REMEMBER

The steepest part of the curve is the part you are present for. Dietary counselling is essential, and advice may be more effective when it comes before weight gain starts, which is the argument for placing it at initiation.

15.4 Who gains, and why the group mean misleads

Almost all antipsychotics have been associated with weight gain and the mean varies substantially between them, but variation between individuals on the same drug is greater still, some losing weight and some gaining a great deal. The source supplying the ranking states plainly that a drug's reported mean may not usefully predict what an individual will gain. Some predictors are consistent. Gain may be more pronounced in antipsychotic-naive patients and early in treating a psychotic illness, and women may be at greater risk than men.

- Weight gain: younger age, female gender, nonwhite race or ethnicity, family history.
- Insulin resistance: a diagnosis of schizophrenia itself, nonwhite race or ethnicity, family history.
- During treatment: rapid weight gain and a rise in plasma triglycerides both predict later diabetes.

The third turns two routine measurements into an early-warning system before any glucose abnormality declares itself. One intuitive lever does not work: within the dose range commonly used for schizophrenia spectrum disorders there is no evidence of a dose relationship for metabolic disorders, so dose reduction is not a strategy to treat them. That is stated for clozapine and is the general rule; the quetiapine finding below is a documented exception.

PITFALL

The reflex response to weight gain on clozapine is to shave the dose. It buys nothing metabolically and risks what clozapine was prescribed to do. The specific interventions belong with the management of antipsychotic metabolic disturbance rather than here.

15.5 The comparative ranking, and the discrepancy inside it

The comparative metabolic risk ranking is best held as ordinal tiers, not a league table of kilograms. The standard rendering clusters antipsychotics into three groups derived from meta-analyses.

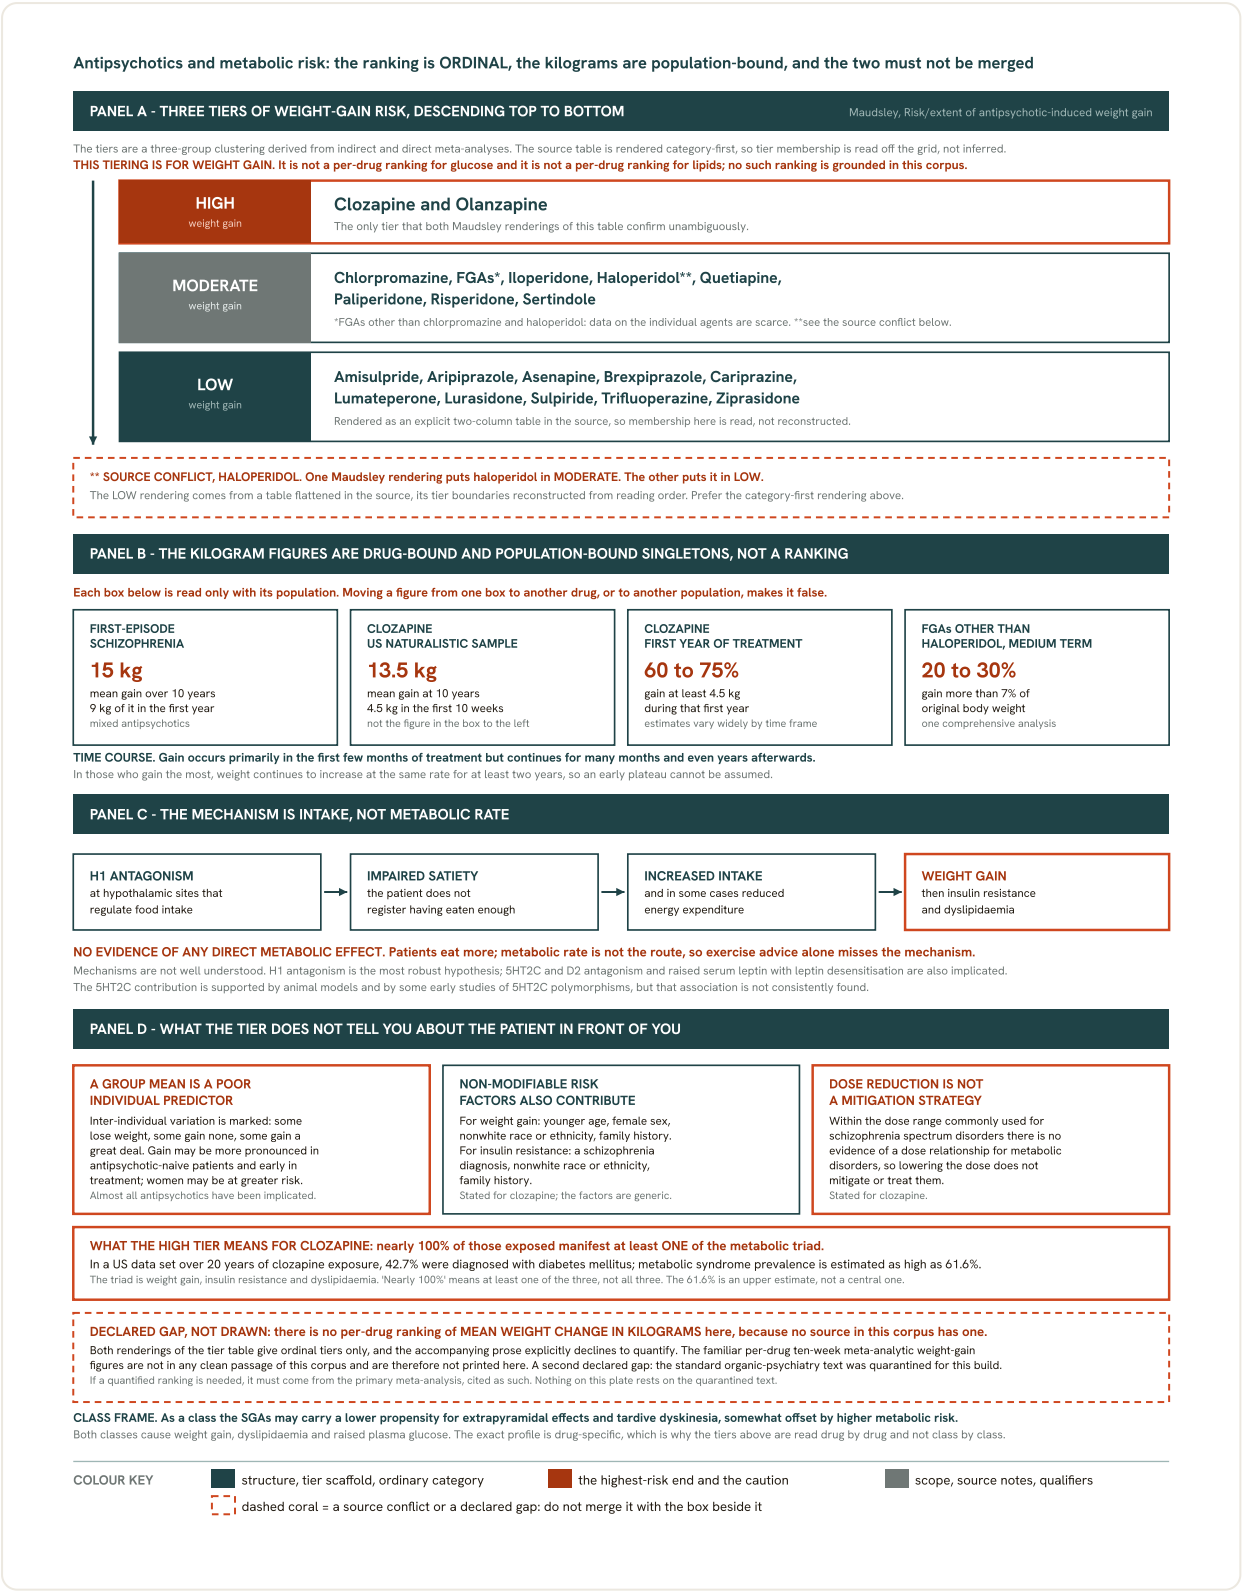


Fig 33.6 Comparative metabolic risk of antipsychotics: an ordinal three-tier ranking for weight gain, with the grounded kilogram figures kept as population-bound singletons. The teaching point is the shape of the evidence, not a league table. The corpus supports High, Moderate and Low tiers for weight gain and nothing finer, and the two available renderings of the same tier table disagree about haloperidol. Every kilogram and every percentage on the plate belongs to a named drug in a named population and does not travel. The mechanism is impaired satiety and increased intake, not a direct effect on metabolic rate, and dose reduction is not a mitigation. (The Maudsley Prescribing Guidelines and Maudsley 15th edition; Stahl's The Clozapine Handbook.)

TIER	AGENTS	READING
High risk or extent of weight gain	Clozapine and olanzapine.	Stable across both renderings and across the lipid and glucose evidence.
Moderate	Chlorpromazine, first-generation antipsychotics as a group, iloperidone, haloperidol, quetiapine, paliperidone, risperidone, sertindole.	The first-generation entry carries the footnote excluding haloperidol.
Low	Amisulpride, aripiprazole, asenapine, brexpiprazole, cariprazine, lumateperone, lurasidone, sulpiride, trifluoperazine, ziprasidone.	Sulpiride and trifluoperazine sit here: the generation label does not predict metabolic risk.

The accompanying figure sets the same ordering out across weight, glucose and lipids. Read the tiers as tiers: no source here ranks antipsychotics by mean weight change in kilograms, and the source prose explicitly declines to quantify. A per-drug kilogram league table asserts what the evidence does not support.

MNEMONIC

Same two drugs, all three lists

Clozapine and olanzapine head the weight-gain tier, head the lipid ranking, and carry the strongest links to impaired glucose tolerance, diabetes and ketoacidosis. The top of all three lists is the same pair.

Now the part that separates a good answer from a safe one. Two renderings of the same three-tier table place haloperidol differently. The fifteenth edition of the Maudsley guidance puts it in the moderate tier, as above; the other Maudsley rendering puts haloperidol in the low tier, with the aripiprazole and ziprasidone group. This is unresolved: it may be an editorial revision, or an artefact of the second table's layout, whose tier boundaries must be reconstructed from a flattened run of drug names, not read off a grid. Consistent across both is that the general first-generation weight-gain statement excludes haloperidol.

TRAP **Being asked to place haloperidol in a weight-gain tier.** Candidates answer confidently in one direction and are marked against the other, because two renderings disagree and neither can be dismissed. Answer what is certain in both: clozapine and olanzapine are the high-risk pair, and haloperidol is not in the high-risk tier. Then name the discrepancy rather than resolving it by assertion.

15.6 Dyslipidaemia: triglycerides are the headline

Among older drugs the phenothiazines raise triglycerides and low-density lipoprotein cholesterol and lower high-density lipoprotein cholesterol, though the magnitude is poorly quantified, while haloperidol appears to have minimal effect on lipid profiles. Cholesterol can rise, but the most profound effect of this class is on triglycerides.

The lipid ranking parallels the weight ranking without duplicating it. Clozapine and olanzapine have the greatest propensity to raise lipids; quetiapine and risperidone are moderate; aripiprazole, lurasidone and ziprasidone have minimal effect and may modestly reverse dyslipidaemia from a previous antipsychotic; cariprazine and brexpiprazole also appear limited. Haloperidol therefore sits at the minimal end for lipids and in the moderate weight tier. That axis-dependence is what a viva turns on.

Olanzapine raises triglyceride levels by **40% over the short term at 12 weeks and the medium term at 16 months**, and levels may keep rising for up to a year. Up to two-thirds of olanzapine-treated patients have raised triglycerides and just under one in ten may develop severe hypertriglyceridaemia. A United Kingdom case-control study found olanzapine-treated patients five times more likely to develop hyperlipidaemia than those with no antipsychotic exposure and three times more likely than those on first-generation drugs; those multipliers have different comparators and are not interchangeable. Risperidone showed no increase against either. Clozapine gives the same picture over a longer horizon, with mean triglycerides doubling and cholesterol rising by at least 10% after 5 years. Both drugs warrant particular care in patients already obese, diabetic or hyperlipidaemic.

The threshold that converts a laboratory abnormality into acute medical risk is **severe hypertriglyceridaemia**, a fasting level above **5 mmol/L**, which is a risk factor for pancreatitis. The sample must be fasting, and this threshold sits far above the much lower triglyceride cut-off used as a metabolic syndrome criterion. Confusing the two produces false alarm or false reassurance.

15.7 Glucose: insulin resistance, diabetes and ketoacidosis

Diabetes is two to three times more prevalent in schizophrenia than population norms, and the observation predates antipsychotics. Drug-attributable risk sits on top of illness-attributable risk, and the two cannot be separated at the bedside.

Clozapine is strongly linked to hyperglycaemia, **impaired glucose tolerance and diabetic ketoacidosis**, with a risk apparently higher than with other second-generation or first-generation drugs, especially in younger patients, though not consistently. As many as a third of patients might develop diabetes after 5 years of clozapine treatment. Timing is wide: many cases in the first six months, some within a month, some only after years. Death from ketoacidosis has been reported, and clozapine-associated diabetes is not necessarily linked to obesity or family history, though both raise risk. Olanzapine mirrors this, and is probably more diabetogenic than risperidone.

Quetiapine and risperidone have both been linked to new-onset diabetes and ketoacidosis, with far fewer reports. Quetiapine appears more likely than first-generation drugs to be associated with diabetes, and two studies found it equal to olanzapine. Here is the one documented dose relationship in this territory: in patients taking quetiapine for a psychotic illness, glycaemic risk may be dose-related, with **400 mg or more per day** clearly linked to changes in glycosylated haemoglobin. That is a metabolic-risk threshold, not a dosing recommendation.

At the favourable end, amisulpride appears not to elevate plasma glucose or be associated with diabetes, and neither aripiprazole nor ziprasidone alters glucose homeostasis, with a large case-control study confirming no increased diabetes risk for amisulpride or aripiprazole. Those three are recommended for patients predisposed to diabetes, or as an alternative to a known diabetogenic drug; lurasidone and asenapine appear to have no effect, and initial data suggest the same for brexpiprazole and cariprazine. One caveat: ketoacidosis has been reported with aripiprazole, so low risk is not no risk. Because between-drug differences are small and inconsistently shown, they may be unimportant in practice: risk is probably increased for everyone with schizophrenia receiving any antipsychotic.

TRAP **Applying the general-population age gradient for diabetes.** Candidates reason that risk climbs with age and name the elderly patient. It runs the other way here. Risk is increased to a much greater extent in younger adults than in the elderly, in whom antipsychotics may show no increased risk at all, and first-episode patients are particularly prone across a variety of agents.

15.8 Detecting it: which test, which threshold, what it means

LOCUS determines what is realistically obtainable. In an outpatient clinic a fasting sample can be arranged for a return visit; on an acute ward or with any disorganised patient, fasting tests are often unobtainable, and insisting means obtaining nothing. In every locus the pre-treatment measurement has the most leverage, separating an abnormality the drug produced from one the patient already had.

FOCUS shifts with phase. Before treatment the target is undiagnosed disease: although ketoacidosis shortly after starting clozapine was an early concern, later studies showed these patients had untreated diabetes, with a mean glycosylated haemoglobin of 13.3% at admission for ketoacidosis. That is a mean at presentation, not a threshold or target. Acutely the focus is symptomatic hyperglycaemia and ketoacidosis; over the short and medium term the rate of weight change and the triglyceride trajectory; long term the accumulated cardiovascular risk.

MODUS is the choice of test, governed by the Maudsley prescribing guidance on antipsychotic monitoring, the American Psychiatric Association practice guideline, and the American Diabetes Association criteria that both reproduce. The oral glucose tolerance test is the most sensitive method and the ideal; fasting plasma glucose is less sensitive but recommended, and any abnormal fasting value should provoke an oral glucose tolerance test. Where fasting is impractical, random plasma glucose or **glycosylated haemoglobin** may be used. Patients should report fatigue, candida infection, thirst and polyuria. The cadence is set out where this chapter deals with metabolic monitoring protocols.

The thresholds are the general diabetes ones. Diabetes is diagnosed by a fasting plasma glucose of at least 126 mg/dl or 7.0 mmol/l, with fasting defined as no caloric intake for at least 8 hours; or a two-hour value of at least 200 mg/dl during a 75 gram oral glucose tolerance test; or a random plasma glucose of at least 200 mg/dl with classic hyperglycaemic symptoms or crisis; or a glycosylated haemoglobin of at least 6.5%, confirmed on repeat testing where hyperglycaemia is

not unequivocal. The American Psychiatric Association states it as higher than 125 mg/dL: the same cut-off expressed differently, not a competing standard.

Interpreting glycosylated haemoglobin means holding its window in mind. It reflects mean blood glucose over the preceding 8 to 12 weeks; below 5.7% is normal, 5.7 to 6.4% indicates increased risk of developing diabetes, 6.5% or above indicates diabetes mellitus, and 7% or less is a common treatment target in established diabetes. That lookback is why the test cannot exclude recent antipsychotic-induced hyperglycaemia: a patient three weeks into olanzapine with a normal result has been told about the two months before the drug started. Metabolic syndrome, meanwhile, is measured against jurisdiction-dependent criteria. The International Diabetes Federation definition uses **ethnicity-specific waist circumference thresholds** whose Asian values apply to Asian-Indian patients, while the definition used for the United States and Canada sets a considerably higher waist threshold. Both require at least three of five risk factors.

IN INDIA

Two divergences bite. Glycosylated haemoglobin is not a universal substitute here: where a haemoglobinopathy or any cause of increased red cell turnover is present, including second-trimester and third-trimester pregnancy, haemodialysis, recent blood loss, transfusion and erythropoietin therapy, fasting blood glucose should be used rather than glycosylated haemoglobin. Haemoglobinopathies are not rare here, so the convenient test is unreliable in exactly the patients from whom a fasting sample is hardest to obtain. Second, Indian consultation-liaison guidance reaches the international conclusion independently: olanzapine and clozapine carry the maximal risk of precipitating metabolic syndrome and should be used with proper clinical prudence in conditions such as diabetes, stroke and heart disease, while aripiprazole may be used as a first-line agent in cardiac disease given its safer metabolic and cardiac profile.

Taken whole, the shape is simple: the same pair at the top of every axis, a broad unremarkable middle, a low-risk group of both newer and older agents, and one drug whose position is contested. Choosing between them belongs to the Schizophrenia: management, antipsychotics and adverse effects notes; this chapter asks only that you know what you are trading away.

Clozapine and olanzapine top the weight, lipid and glucose rankings alike; the three effects dissociate, so weight, a fasting lipid panel and a glucose measure are three separate obligations, and dose reduction is a remedy for none of them.

16 · Metabolic monitoring protocols and parameters for patients on antipsychotics

This is the part of antipsychotic prescribing every guideline calls mandatory and almost nobody does properly. The question always has three parts: which parameters you record, at what intervals you repeat them, and what you do with a value outside the reference range. The first is nearly uncontested and the third is where marks are won, because a schedule generating numbers nobody acts on is an audit exercise rather than care. The second is where candidates lose marks, because the major guidelines give different intervals for the same test in the same patient.

Four bodies of guidance govern this territory and each is internally coherent. The Maudsley Prescribing Guidelines carry a parameter-by-parameter table with drug-specific intensification and an action column. NICE guidance, reproduced in that chapter, specifies a baseline panel. The American Psychiatric Association practice guideline gives a separate schedule and is candid that there are no absolute criteria for frequency, so decisions must follow clinical circumstance. Stahl's The Clozapine Handbook gives a clozapine-specific routine synthesised from published recommendations rather than issued by one body. Kaplan and Sadock's Comprehensive Textbook of Psychiatry reproduces the older multi-society consensus table and gives its own prose recommendations, and the two disagree. Hold them as four positions, not one consensus.

8% TO UNDER 30% PATIENTS GETTING BASELINE GLUCOSE AND LIPID CHECKS	**8.8%** GETTING ANY FOLLOW-UP ASSESSMENT	**23.0 kg/m2** OVERWEIGHT THRESHOLD, ASIAN CUTPOINTS	**1 BMI UNIT** RISE THAT SHOULD TRIGGER INTERVENTION

16.1 The baseline panel: what is recorded before the first dose

Everything downstream depends on a pre-treatment measurement, because without one you cannot separate an abnormality the drug produced from one the patient brought with them. A raised fasting glucose before the first tablet is a comorbidity to treat; the same value four months in is a reason to reconsider the drug. NICE sets out the **baseline investigations** to be undertaken and recorded before starting.

- Weight.
- Waist circumference.
- Pulse and blood pressure.
- Fasting blood glucose, HbA1c, a blood lipid profile and prolactin.
- Assessment of movement disorders.
- Assessment of nutritional status, diet and level of physical activity.

Note that the panel asks for both a fasting glucose and a glycated haemoglobin in everyone, and for a prolactin level in everyone. An electrocardiogram is added selectively, where the summary of product characteristics specifies it, where examination has identified a specific cardiovascular risk such as a diagnosis of high blood pressure, where there is a personal history of cardiovascular disease, or where the person is being admitted as an inpatient. The American baseline panel is narrower and differently shaped: body weight, height and BMI; screening for diabetes risk factors with a fasting blood glucose; a lipid panel; a determination of whether metabolic syndrome criteria are met; pulse and blood pressure; and screening for symptoms of hyperprolactinaemia, with a prolactin level only if the clinical history indicates one. The criteria that determination uses are set out where this chapter deals with metabolic syndrome, obesity and mortality in serious mental illness.

MNEMONIC

Weight, Waist, Pressure, Bloods, Movements, Lifestyle

The six headings of the baseline panel in the order they are recorded. Five of the six need no laboratory.

16.2 The architecture of follow-up, and where the sources part company

Every schedule shares one skeleton: a baseline, intensified early surveillance, a consolidating check inside the first six months, then a maintenance interval. What differs is the length of the intensified period, the timing of the consolidating check and, most consequentially, the maintenance floor. The accompanying figure lays that skeleton out.

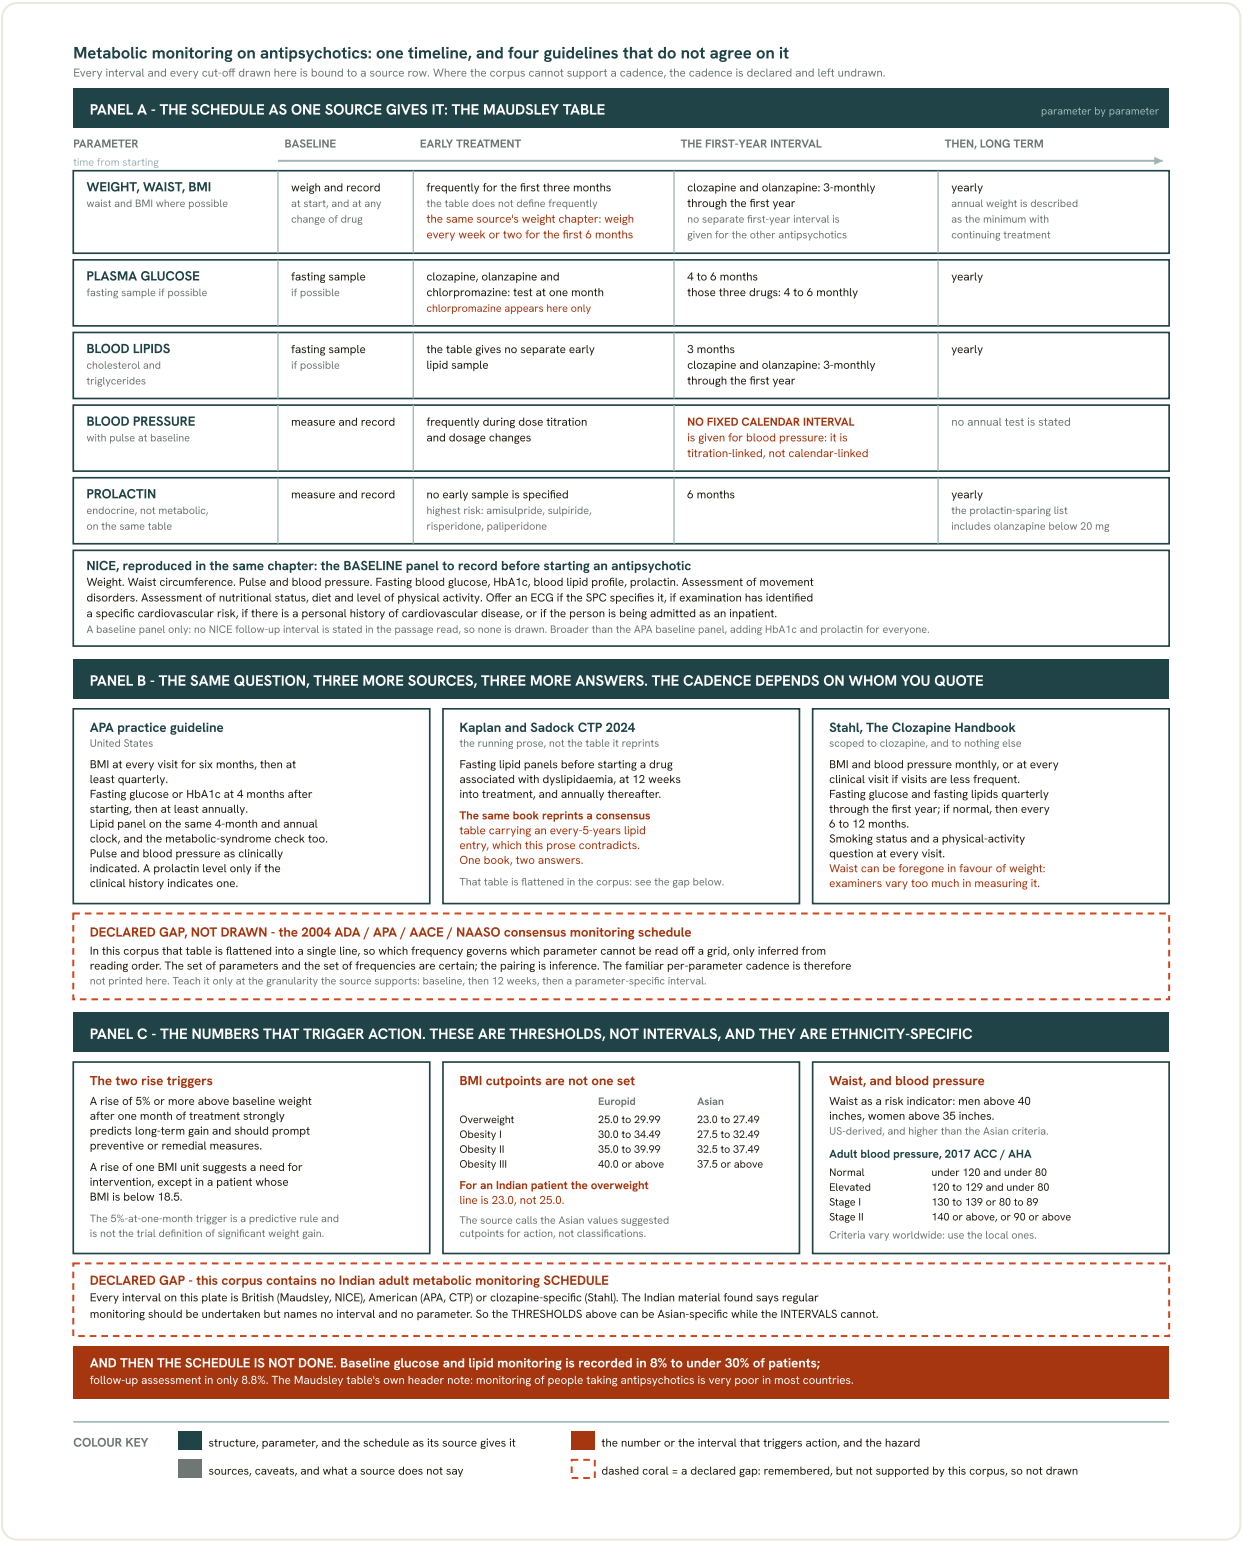


Fig 33.7 Metabolic monitoring on antipsychotics: one timeline, and four guidelines that do not agree on it. The Maudsley table is the only complete parameter-by-parameter schedule in this corpus, and its blood-pressure row is titration-linked rather than calendar-linked, so there is no annual blood-pressure interval to quote. Chlorpromazine joins clozapine and olanzapine in the special-precaution list for glucose and for glucose only. The APA guideline, the CTP prose and the clozapine-specific Stahl schedule each answer the same question differently, so the cadence you quote depends on whom you quote. The anthropometric thresholds can be Asian-specific, but every interval drawn here is British or American: no Indian adult monitoring schedule is present in this corpus, and the flattened 2004 consensus table is declared and left undrawn rather than reconstructed from memory. (The Maudsley Prescribing Guidelines, and NICE as reproduced there; the American Psychiatric Association practice guideline; Kaplan and Sadock's Comprehensive Textbook of Psychiatry 2024; Stahl, The Clozapine Handbook.)

PARAMETER	THE MAUDSLEY PRESCRIBING GUIDELINES	APA PRACTICE GUIDELINE	STAHL'S THE CLOZAPINE HANDBOOK (CLOZAPINE)
Weight and BMI	Baseline, frequently for three months, then yearly; clozapine and olanzapine, 3-monthly for the first year, then yearly	BMI every visit for 6 months, then at least quarterly	Monthly, or every clinical visit if visits are less frequent
Blood pressure	Baseline, then frequently during titration and dose changes; no fixed calendar interval	Pulse, blood pressure and temperature as clinically indicated	Monthly, or every clinical visit if less frequent
Fasting glucose	Baseline, at 4 to 6 months, then yearly; clozapine, olanzapine and chlorpromazine at baseline, one month, then 4 to 6 monthly	Fasting glucose or HbA1c at 4 months, then at least annually	Quarterly in the first year, then every 6 to 12 months if normal
Fasting lipids	Baseline, at 3 months, then yearly; clozapine and olanzapine 3-monthly for the first year, then yearly	Lipid panel at 4 months, then at least annually	Quarterly in the first year, then every 6 to 12 months if normal
Prolactin	Baseline, at 6 months, then yearly	Symptom screening each visit until stable, then yearly on a prolactin-raising drug; a level only if the history indicates one	Not among this routine table's parameters

An older multi-society consensus sits behind much of this. The **2004 consensus development conference on antipsychotic drugs, obesity and diabetes**, convened by the American Diabetes Association, the American Psychiatric Association, the American Association of Clinical Endocrinologists and the North American Association for the Study of Obesity, produced the table most later schedules descend from: personal and family history, weight and BMI, waist circumference, blood pressure, fasting plasma glucose and fasting lipid profile, on a baseline and 12-week frame. Reproduce it only at the granularity you can defend: in the rendering available here the pairing of interval to parameter is reading order rather than a grid, and the consensus called these routine recommendations that specific situations may warrant intensifying.

16.3 Weight: the parameter that moves first and costs nothing

Body weight is the most informative single measurement here because it changes earliest, needs no laboratory and predicts the rest. The minimum standard is that every patient starting or changing antipsychotic medication is weighed and the weight clearly recorded, ideally with BMI and waist circumference. The intensity that matters is early: weighing every week or two for at least the first six months, more often than any parameter table above asks for.

That frequency exists to catch one signal. **Rapid weight gain in early treatment**, taken as an increase of **at least 5% above baseline after a month of treatment**, strongly predicts long-term weight gain and should prompt consideration of preventative or remedial measures. It converts a month of observation into a prediction about years, long before any laboratory value moves. With continuing treatment, annual weighing is the stated minimum, and that floor is where guidelines diverge.

The second trigger is expressed in BMI. A BMI of 25 to 29.9 is overweight and 30 or above is obese, and, except where BMI is below 18.5, a rise of **1 unit** should prompt intervention: closer weight monitoring, a weight management programme, an adjunctive treatment to reduce weight, or a change of antipsychotic. The low-BMI exemption is dropped more often than remembered, and it matters: an underweight patient gaining a unit of BMI is doing something desirable. What to do once the trigger fires belongs where this chapter deals with managing antipsychotic metabolic disturbance and hyperprolactinaemia; drug selection belongs to the Schizophrenia: management, antipsychotics and adverse effects notes.

TRAP **Confusing the early-prediction percentage with the trial definition of clinically significant weight gain.** The figure triggering action at one month is a predictive threshold, set low so it fires early. The percentage of baseline weight used in trials to define clinically significant gain is a different and higher figure, taught with the antipsychotic metabolic effects material. Two percentages, two jobs; quote the wrong one and the answer changes from surveillance to a switching decision.

16.4 Anthropometric thresholds, and why the Indian ones are different

Waist circumference as a risk indicator is where a Western protocol becomes actively misleading in an Indian clinic. The general health-risk figures in the standard textbook are that men with a waist larger than 40 inches and women with a waist larger than 35 inches are at increased health risk. Those are American-derived. The International Diabetes Federation criteria, ethnicity-specific and with materially lower Asian values, are set out where this chapter deals with metabolic syndrome, obesity and mortality in serious mental illness. The same problem recurs with BMI.

CATEGORY	EUROPID INDIVIDUALS	ASIAN INDIVIDUALS
Overweight	25.0 to 29.99 kg/m2	23.0 to 27.49 kg/m2
Class I obesity	30.0 to 34.49 kg/m2	27.5 to 32.49 kg/m2
Class II obesity	35.0 to 39.99 kg/m2	32.5 to 37.49 kg/m2
Class III obesity	40.0 kg/m2 or above	37.5 kg/m2 or above

These are offered as **suggested cutpoints for action** rather than as a reclassification of obesity, a hedge worth reproducing. A methodological dissent follows: the clozapine source recommends foregoing waist circumference for the more accurate weight, given significant variation between examiners, whereas the Maudsley table and NICE both keep monitoring waist circumference in the protocol.

IN INDIA

The thresholds are Indian, the intervals are not. Ethnicity-specific cutpoints apply directly here, with overweight beginning at a BMI of 23.0 kg/m2 rather than 25.0, so a patient carrying what a Western chart calls a normal BMI can already be in the band where action is suggested. Blood pressure criteria carry the same warning: the source cautions that they are the American criteria as modified in 2017, that criteria vary worldwide, and that clinicians should know those used locally. What no source here supplies is an Indian adult schedule for antipsychotic metabolic monitoring; Indian consultation-liaison guidance says regular monitoring should be undertaken but specifies neither parameters nor intervals. Take thresholds from the ethnicity-specific tables and timing from Western guidance, and say so rather than imply a domestic schedule.

16.5 Glucose: the interval nobody agrees on

Baseline and follow-up fasting glucose appears in every schedule above, and the parameter is a fasting plasma glucose where a fasting sample can be had; an abnormal value draws a graded response: lifestyle advice, a fasting sample or a non-fasting sample with a glycated haemoglobin, and referral to the general practitioner or a specialist. Glucose carries the most explicit drug-specific intensification. Clozapine, olanzapine and chlorpromazine attract a check at **one month** that other antipsychotics do not; note which drug joins the usual pair, since chlorpromazine is on the special-precaution list for glucose but not for lipids or weight.

The routine interval has three answers. The Maudsley table rechecks at four to six months and then yearly. The American guideline rechecks a fasting glucose or a glycated haemoglobin at four months and then at least annually, with no drug-specific intensification. The clozapine handbook tests quarterly through the first year and then every six to twelve months if normal, adding that monthly or bimonthly fasting glucose should be considered in the first six months where multiple diabetes risk factors are present, and adding a glycated haemoglobin where fasting values are hard to obtain. Monitoring for **impaired fasting glucose** is recommended whatever the drug, because prevalence in this population is high enough that nobody is exempt.

TRAP **Answering "the" recheck interval as though one exists.** Four sources give four lipid schedules and three give three for glucose, and the Maudsley guidance concedes in its own prose that there is no clear consensus on diabetes monitoring for antipsychotic-treated patients and that formal guidelines vary considerably. The safe answer names the source with the interval, notes the divergence, then gives the item most often missed: the one-month glucose check for clozapine, olanzapine and chlorpromazine.

16.6 Lipids: four schedules and a source that contradicts itself

A fasting sample is preferred where obtainable, and cholesterol and triglycerides are the reported values. The action column is more interventionist than for glucose: lifestyle advice, and consideration of both a change of antipsychotic and a statin. Some agents, aripiprazole and lurasidone among them, are not clearly associated with dyslipidaemia, yet everyone is monitored, because prevalence in this population is high. Population risk rather than drug risk sets the monitoring floor across all three parameter tables.

The intervals are the most divergent set in the topic. The Maudsley table rechecks at 3 months, then yearly, with clozapine and olanzapine 3-monthly through the first year. The American guideline rechecks at four months and then at least annually. Kaplan and Sadock's prose evaluates fasting lipid panels before starting a medication associated with dyslipidaemia, at 12 weeks into treatment, and annually thereafter. The clozapine handbook runs quarterly through the first year and then six to twelve monthly. Three months, four months and twelve weeks make no clinical difference and are examinable separately.

PITFALL

The outlier deserves naming rather than quiet omission. The reproduction of the 2004 consensus table sets the fasting lipid profile at baseline, 12 weeks and then every five years, a maintenance interval the same textbook's own prose contradicts by specifying annually. Quoting a five-yearly lipid check as current practice quotes a position later guidance moved away from. Present it as history, not protocol.

16.7 Blood pressure and prolactin: the two parameters that are not metabolic

Two parameters on the monitoring sheet answer to a different logic. Blood pressure is not calendar-linked in the Maudsley scheme at all: baseline, then frequently during titration and dose changes, because the risk it tracks is haemodynamic and dose-driven rather than cumulative. Clozapine, chlorpromazine and quetiapine are most likely to produce postural hypotension, while amisulpride, aripiprazole, brexpiprazole, cariprazine, lurasidone, trifluoperazine and sulpiride are listed as not requiring the same precaution. Slow titration where severe hypotension, or hypertension on clozapine, is seen; consider switching to another antipsychotic where postural hypotension is symptomatic; treat hypertension along standard guideline lines.

Staging the reading matters because the thresholds moved. Against the **2017 American College of Cardiology and American Heart Association criteria**, normal is a systolic below 120 mmHg and a diastolic below 80 mmHg; elevated is a systolic of 120 to 129 mmHg with a diastolic below 80 mmHg; stage I hypertension is a systolic of 130 to 139 mmHg or a diastolic of 80 to 89 mmHg; and stage II is a systolic of 140 mmHg or more or a diastolic of 90 mmHg or more. Note the connectives: normal and elevated require both criteria, the hypertensive stages either.

Prolactin is measured at baseline, at six months and then yearly, and the drugs attracting attention are amisulpride, sulpiride, risperidone and paliperidone. Several usually do not raise it, including asenapine, aripiprazole, brexpiprazole, cariprazine, clozapine, lurasidone, quetiapine, ziprasidone and, with a dose qualification easy to lose, olanzapine at **less than 20 mg**. Where hyperprolactinaemia is confirmed and symptomatic the response is to switch, and **bone mineral density testing such as DEXA scanning** should be considered where levels are chronically raised. Management belongs where this chapter deals with managing antipsychotic metabolic disturbance and hyperprolactinaemia.

RULE **Every parameter carries an action, and the action is what makes it monitoring rather than measurement.** Each row of the standard tables pairs a test with a response: lifestyle advice, a confirmatory test, a drug change, a pharmacological intervention, or a referral. If you cannot say what you would do with an abnormal value before ordering it, you are collecting data, not monitoring.

16.8 Making the protocol actually happen

LOCUS decides what is obtainable. On an inpatient ward the fasting sample, weight and blood pressure are all easy and the baseline panel should be completed before the first dose; admission itself triggers a baseline electrocardiogram. In the outpatient clinic the fasting sample is the constraint, which is why the alternatives are written into the guidance: a non-fasting sample with a glycated haemoglobin. In shared care, abnormal values are referred onward, making the referral pathway part of the protocol rather than an afterthought.

FOCUS shifts by phase. Before the first dose the target is undiagnosed existing disease and a clean set of reference values. Through the first weeks it is the rate of change in weight, the only signal moving fast enough to inform this early. Between three and six months it moves to the laboratory parameters as the consolidating checks fall due. In maintenance it is accumulating cardiovascular burden, and here the sources disagree most sharply: annual, at least quarterly, or monthly for clozapine.

MODUS is the set of instruments, and most are not tests. Scales, a tape measure and a sphygmomanometer cover four parameters. Two further items belong on the clozapine routine sheet and are useful for everyone: smoking status, asked and quantified at every visit, with cessation programmes and medication offered quarterly once the patient is psychiatrically stable, and a **physical activity vital sign**, asked and recorded at every monthly visit. Framing activity as a vital sign rather than as advice is the point.

Build the metabolic monitoring protocol into the structure of the clinic rather than into individual intention, because intention demonstrably fails. Screening for metabolic risk occurs at very low rates in schizophrenia even where the prescribed antipsychotic is one associated with glucose and lipid disturbance, and despite recommendations from the American Psychiatric Association, the American Diabetes Association and other groups, and despite regulatory warnings. The Maudsley table prefaces itself in plainer words: monitoring of people taking antipsychotics is very poor in most countries.

GOOD TO REMEMBER

Put the parameters on the prescription sheet, so the review that renews the drug cannot be completed without them.

Baseline everything before the first dose, weigh every week or two through the early months and act on a 5% rise at one month or a 1-unit rise in BMI, recheck glucose and lipids between three and six months and then at your guideline's maintenance floor, and remember the extra one-month glucose check for clozapine, olanzapine and chlorpromazine.

17 · Managing antipsychotic metabolic disturbance and hyperprolactinaemia

Knowing that antipsychotics wreck metabolism and raise prolactin earns nothing on its own. The marks are for the sequence: what you try first, what you switch to, what you refuse to add, and when a metabolic complication stops the drug. Both problems are iatrogenic, and both are found by a measurement rather than reported by the patient.

Detection aside, drug selection belongs to the Schizophrenia: management, antipsychotics and adverse effects notes; metabolic syndrome criteria to where this chapter deals with metabolic syndrome, obesity and mortality in serious mental illness; the measurement schedule to where it deals with metabolic monitoring protocols; the reproductive-endocrine view of a raised prolactin to where it deals with reproductive hormones and mood.

500-2000 MG/DAY METFORMIN, ANTIPSYCHOTIC WEIGHT GAIN	**2.5-3 G/DAY** METFORMIN, PROLACTIN LOWERING	**3 MG/DAY** LIRAGLUTIDE, WEIGHT LOSS INDICATION	**UNDER 25** AVOID HIGH-PROLACTIN DRUGS, PRE-PEAK BONE MASS

17.1 The order of moves for a patient who has gained weight

LOCUS. Almost all outpatient: what works is what a patient sustains between appointments. Admission supervises a structured diet and activity programme and gives a switch daily observation. Emergency presentation belongs to the acute metabolic decompensations, not to weight.

FOCUS is phase-dependent: acutely, nothing to treat but the recognition; short term, the rate of gain, steepest early; medium term, reversal of accrued gain with its lipid and glucose consequences; long term, cardiovascular risk and willingness to stay on an effective antipsychotic at all.

MODUS. The Maudsley prescribing guidance sets the sequence, the American Psychiatric Association practice guideline agreeing in spirit. **Behavioural lifestyle programmes** improving diet and physical activity come first, no token gesture: meta-analyses show a robust effect for prevention and treatment alike, and early intervention prevents gain rather than chasing it. Alongside or instead, switching antipsychotic to reduce metabolic risk is the other first-line option, with fairly strong support for aripiprazole, ziprasidone or lurasidone for reversing gain already accrued, though it risks relapse and discontinuation. **Adjunctive aripiprazole** is the related option, weight loss observed when added to clozapine or olanzapine. The effect of both is modest and many patients remain overweight.

Pharmacological treatment sits after those two, only where behavioural strategies or a switch have failed, or where obesity already poses an immediate physical risk. **Metformin** is then probably the drug of choice for prevention and treatment alike, though glucagon-like peptide-1 agonists may ultimately prove better. Bariatric surgery may have a place in the rare severe case where all else has failed, its efficacy for drug-induced gain unknown.

MNEMONIC

Lifestyle, Switch, Metformin, then the rest

The order the Maudsley guidance puts it in, the rest being thinner evidence. Clozapine is the documented exception to the ordering, not to the content.

17.2 Metformin: the exception that reverses the order

Metformin for antipsychotic-induced weight gain has the largest evidence base of any agent here. In non-diabetic patients it reduces and reverses the gain at 500 to 2000 mg per day, with benefit on other metabolic parameters, the evidence mainly for olanzapine-induced gain. Individual studies are not unanimous, but meta-analyses show a clear effect, with a positive randomised trial and extension study in children and adolescents with autism spectrum disorder. Two editions of the same guidance give the identical 500 to 2000 mg per day range for antipsychotic-induced weight gain; the second adds that a Cochrane review found only low-certainty evidence that metformin prevents gain in schizophrenia, and that a later cohort study found it did prevent gain with clozapine. It suits weight gain coexisting with diabetes or polycystic ovary syndrome.

Against that ordering stands clozapine-specific guidance. Unless renal function forbids it, every patient starting clozapine should commence metformin concurrently: gain with clozapine is pervasive and starts early, and metformin brings benefits beyond weight, including triglyceride reduction and delayed progression to type 2 diabetes in prediabetic patients. Prophylaxis at initiation, not rescue after failure.

TRAP **Being asked when metformin for antipsychotic-induced weight gain should be started, and answering with one rule.** Two clean sources sit in opposite positions: the Maudsley guidance puts drug treatment last, after behavioural strategies and switching have failed; the clozapine-specific guidance starts metformin on the same day as clozapine, for everyone. State the partial reconciliation rather than assuming it: the prophylactic stance is explicitly clozapine-specific, and switching is largely unavailable on clozapine anyway. Present neither as the universal rule.

17.3 Starting and titrating metformin without losing the patient

For clozapine-associated weight gain and metabolic adverse effects, begin at 500 mg orally each morning, since higher initial doses cause diarrhoea and nausea; extended-release forms may help further. Titrate over three weeks: 500 mg each morning in the first week, then 500 mg twice daily or 1000 mg of the extended-release preparation each morning in the second, then 1000 mg twice daily or 2000 mg extended-release each morning from the third week if tolerated. If 2000 mg per day is not tolerated, fall back to 1500 mg per day, as 750 mg twice daily or 1500 mg extended-release each morning.

Renal function governs eligibility. Metformin for antipsychotic-induced weight gain is safe at an estimated glomerular filtration rate of 30 ml/min or above, with use reviewed below 45 ml/min. The ladder: yearly eGFR while it is 60 ml/min or above; every 3 to 6 months once below 60; below 45, halve the metformin dose and check eGFR every 3 months; stop metformin altogether below **30 ml/min**. Check vitamin B12 yearly: deficiency is commoner on metformin, and its cognitive and affective consequences are readily misread as illness rather than as iatrogenic and correctable.

GOOD TO REMEMBER

Metformin is usually abandoned for one of two reasons: a gastrointestinal start that was too fast, or a renal result nobody looked at. Both are procedural. Write the titration on the prescription, default to extended-release where cost allows, and put eGFR and vitamin B12 on the annual bloods.

17.4 The rest of the pharmacological menu

Four other agents are positively recommended, each with a narrower indication than candidates assume. **Adjunctive aripiprazole** for antipsychotic-induced weight gain is given at 5 to 15 mg per day, randomised trials showing benefit on weight and possibly other metabolic parameters. The scope is strict: weight gain on clozapine or olanzapine only, explicitly not other antipsychotics. It appears safe and unlikely to worsen psychosis. Clozapine-specific guidance restricts it to clozapine monotherapy, and stops it for adverse effects or if no benefit has appeared by 12 to 16 weeks.

Liraglutide, a glucagon-like peptide-1 agonist approved for type 2 diabetes and later as an anti-obesity agent, is given for weight loss at 3 mg per day by subcutaneous injection, higher than the diabetes dose of 1.8 mg or less per day; the two must not be collapsed. Data in drug-induced gain are limited, but one randomised trial showed significant weight loss in overweight prediabetic patients stable on olanzapine or clozapine. It is well tolerated apart from gastrointestinal disturbance, and is recommended in prediabetic or diabetic patients and for clozapine-induced gain.

Orlistat, at 120 mg three times daily before or after food for medication-induced weight gain, works reliably in obesity with calorie restriction; data in drug-induced gain are few. Two qualifications are easily lost: in clozapine-induced or olanzapine-induced gain the effect appeared only in men, and without calorie restriction the effect in psychiatric patients is very limited. Failure to keep to a low-fat diet produces fatty diarrhoea and possible malabsorption of oral medication, here perhaps the antipsychotic itself. It remains a good choice for clozapine-induced gain, reducing weight and constipation alike. **Topiramate**, at up to 300 mg daily in the same indication, reliably reduces weight even when the gain is medication-induced, with a greater effect for prevention than treatment. Its problem is neuropsychiatric: sedation, confusion and cognitive impairment.

17.5 What the evidence does not support

An examiner can as easily ask what you would not prescribe. The agents below were assessed for antipsychotic-induced weight gain and are not recommended; the figures are doses studied, never recommendations, and several still appear on ward rounds on a plausible mechanism.

AGENT	DOSE STUDIED IN ANTIPSYCHOTIC-INDUCED WEIGHT GAIN	VERDICT
Amantadine	100-300 mg/day	Too limited to recommend
Alpha-lipoic acid	1200 mg/day	Not recommended
Betahistine	48 mg/day	Limited data; not recommended
Bupropion	no dose recorded	Few data in drug-induced gain; not recommended
Bupropion with naltrexone	no dose recorded	No data in drug-induced gain; not recommended, but not ruled out
Fluvoxamine	50 mg/day	Too limited to recommend, and markedly raises clozapine levels: extreme caution
Melatonin	up to 5 mg at night	Effect, if any, is small
Modafinil	up to 300 mg/day	Not recommended
Naltrexone	25-50 mg/day	Two small pilot randomised trials only
Zonisamide	100-600 mg/day	Not recommended

Agents recommended in earlier editions, including the histamine H2 antagonists, have been removed, though the older textbook a candidate revised from may still carry them.

17.6 Lipids, glucose, and when a metabolic problem stops the drug

Patients with raised cholesterol may benefit from dietary advice, lifestyle change and statins, which seem to be effective in this patient group though interactions are possible. Evidence supports treating cholesterol as low as **4 mmol/l** in high-risk patients, the highest level the National Institute for Health and Care Excellence recommends for secondary prevention of cardiovascular events; no target is set for primary prevention, though recent advice promotes statins wherever ten-year cardiovascular risk exceeds 10%. Separately, coronary heart disease and stroke risk fall by a third if cholesterol is reduced to as low as 3.5 mmol/L. Where triglycerides alone are raised, a low-saturated-fat diet, fish oil and fibrates work without proven mortality benefit, and such patients need screening for impaired glucose tolerance and diabetes. For moderate to severe hyperlipidaemia, switch first to a less offending agent; D2 partial agonists such as aripiprazole are treatments of choice after previous antipsychotic-induced dyslipidaemia.

The glucose lane follows the same logic. Switching to an antipsychotic of lower cardiometabolic risk often reverses changes in glucose tolerance, most compellingly to aripiprazole and also

ziprasidone; for prediabetes or diabetes on clozapine, olanzapine or quetiapine, switch to aripiprazole, brexpiprazole, cariprazine, lurasidone or ziprasidone. Lifestyle intervention prevents diabetes and belongs to everyone with schizophrenia, not only those whose glucose has drifted.

Then comes the question the lane is really testing. Metabolic adverse effects are almost never grounds for stopping clozapine. Weight gain, hypertension, dyslipidaemia and metabolic syndrome are each so classed. The exceptions are emergencies: out-of-control diabetes, diabetic ketoacidosis or a hyperosmolar hyperglycaemic state, during which temporary cessation may help. Developing diabetes is not itself a reason to stop, and no evidence supports dose reduction for glycaemic control. The reflex to shave the dose costs the patient the only drug that worked and buys nothing.

17.7 Hyperprolactinaemia: mechanism, and which drugs actually do it

Dopamine inhibits prolactin release, so dopamine antagonists raise plasma prolactin, probably dose-relatedly. For most antipsychotics the D2 occupancy at which prolactin starts to rise sits very close to that needed for efficacy, leaving little window to exploit. Hence the effect is called probably unavoidable in practice with first-generation drugs, whose main drawbacks are acute extrapyramidal symptoms, hyperprolactinaemia and tardive dyskinesia.

The correction that carries marks: hyperprolactinaemia is usually worse with risperidone, paliperidone and amisulpride than with typical drugs. The generation label predicts nothing; the three-tier ranking makes this explicit. **Prolactin-sparing**, increase very rare: aripiprazole, asenapine, brexpiprazole, cariprazine, clozapine, iloperidone, lumateperone, pimavanserin, quetiapine. **Low-risk prolactin-elevating**, minor changes: lurasidone, olanzapine, ziprasidone. **High-risk prolactin-elevating**, major changes: amisulpride, paliperidone, risperidone, sulpiride and the first-generation antipsychotics. Clozapine and olanzapine head the metabolic lists yet sit in the sparing and low-risk prolactin groups: metabolically brutal may mean endocrinologically gentle, and the reverse.

High-risk prolactin-elevating drugs should be avoided where possible in four groups.

- Patients under 25 years of age, before peak bone mass is reached.
- Patients with osteoporosis.
- Patients with a history of hormone-dependent breast cancer.
- Young women.

The first bears directly on first-episode psychosis, where the typical patient is younger than the threshold and risperidone is a common first choice. The rationale is skeletal, not cosmetic: prolactin-driven hypogonadism during bone accrual costs peak bone mass never fully recovered.

17.8 What the number means, and the three normal ranges

Hyperprolactinaemia is often superficially asymptomatic, and some evidence suggests it does not affect subjective quality of life. The silence misleads. The consequences of hyperprolactinaemia accrue whether or not the patient reports anything, because persistent elevation suppresses the

hypothalamic-pituitary-gonadal axis: sexual dysfunction, menstrual disturbance, breast growth and galactorrhoea, occasionally delusions of pregnancy, and long term reduced bone mineral density and a possible increase in breast cancer risk. That last is hedged in the source and must stay hedged. In premenopausal women, above 100 ng/mL brings infertility and amenorrhoea, hot flushes and vaginal dryness; 50 to 100 ng/mL may cause amenorrhoea or oligomenorrhoea; 20 to 50 ng/mL may cause infertility even with a normal cycle, though most have no galactorrhoea. In men, reduced libido, erectile dysfunction, fatigue, gynaecomastia and infertility, with galactorrhoea only rarely; chronic elevation costs body hair, muscle mass and bone.

Before attributing any of it to the antipsychotic, clear the differential: prolactin also rises with stress, pregnancy and lactation, seizures, renal impairment and other medical conditions including prolactinoma. Venepuncture in a distressed patient is itself a cause.

PITFALL

The raised prolactin is read off the report and blamed on the drug, and a pregnancy that would have changed everything is missed, or a prolactinoma grows behind a pharmacological alibi. Exclude pregnancy, recent seizure, renal impairment and acute stress at sampling first. The corollary error: an elevation on a prolactin-sparing agent dismissed as expected, when nothing about that drug predicts it.

Three sources give three different definitions of normal, and the disagreement is substantive rather than rounding.

SOURCE	NORMAL RANGE	WHAT THE NUMBER TRIGGERS
Maudsley 15th edition	Sex-specific: women 0-25 ng/ml, about 0-530 mIU/L; men 0-20 ng/ml, about 0-424 mIU/L	Elevated is 25-118 ng/ml (530-2500 mIU/L): assess prolactin-related adverse effects systematically and discuss the consequences of prolonged elevation. Highly elevated is above **118 ng/ml, above 2500 mIU/L**: refer to rule out prolactinoma.
Kaplan and Sadock's Comprehensive Textbook of Psychiatry	Not sex-specific: 5 to 20 ng/mL, a lower bound the Maudsley range lacks	Contextual: imaging is not usually performed if the patient is on an antipsychotic known to cause hyperprolactinaemia and the magnitude fits a drug cause
Indian consultation-liaison guidance	Not sex-specific; defined as more than 20 ng/ml	Prompts the search for a contributing drug
Dulcan's Textbook of Child and Adolescent Psychiatry	Not stated	Above 200 ng/mL, or persistent elevation despite a change to a prolactin-sparing drug, image the sella turcica for pituitary adenoma or parasellar tumour; if normal, replace sex steroids to prevent osteoporosis

A woman whose prolactin sits just above the American textbook's upper bound but below the Maudsley ceiling for women is normal by the Maudsley range, above range by that textbook, and hyperprolactinaemic by the Indian definition. Only one sets a different ceiling for women than men, and that sex-specificity is itself the disagreement worth teaching. The imaging thresholds diverge almost twofold, and the child and adolescent figure comes from a different population in a truncated sentence, so neither dismisses the other. Monitoring is contested the same way. The Maudsley fifteenth edition measures plasma prolactin at baseline in every patient, with symptom enquiry at 3 months and a level obtained then only if hyperprolactinaemia is suspected or the patient is on a prolactin-elevating antipsychotic. The same publisher's other edition sets a different routine interval; the American Psychiatric Association guideline obtains a level only where clinical history indicates it; the National Institute for Health and Care Excellence puts prolactin in its universal baseline panel. The cadence belongs where this chapter deals with metabolic monitoring protocols; what belongs here is that a routine baseline level is not universally agreed to be necessary.

17.9 Treating it: what to change, what to add

LOCUS is again the outpatient clinic, switching and adding both carrying a psychiatric risk managed across weeks with the patient at home. **FOCUS** divides cleanly: acutely, whichever symptom brought it to attention, usually sexual dysfunction or amenorrhoea; long term, bone. **MODUS** for antipsychotic-induced hyperprolactinaemia follows the Maudsley guidance, the Indian consultation-liaison position converging on the same manoeuvres.

RULE **Treatment of hyperprolactinaemia depends more on symptoms and long-term risk than on the reported plasma prolactin level.** This is why the guidance stating it offers no action threshold, and it is the safest framing when sources disagree about where normal ends. A markedly raised but asymptomatic level in a patient with no bone risk factors may reasonably be observed after a shared discussion of the consequences of hyperprolactinaemia; a modestly raised level in a young woman with amenorrhoea demands action.

Switching to an antipsychotic of lower prolactin liability is usually the first choice once treatment is judged necessary, though it risks destabilising the illness. Adding aripiprazole is the alternative. Bone-protective measures run in parallel and cost little: stopping smoking, weight-bearing exercise, adequate calcium and vitamin D3. Where a patient must stay on a prolactin-elevating agent and cannot tolerate aripiprazole, **dopamine agonists** work; amantadine, cabergoline and bromocriptine have all been used, each with at least a theoretical potential to worsen psychosis, unreported in trials.

The dose figures for antipsychotic-induced hyperprolactinaemia are where the guidance comes apart. In the management prose, aripiprazole lowers prolactin dose-dependently: 3 mg per day is effective and 6 mg per day more so, **6 mg per day** appearing sufficient. The summary box in the same guidance instead gives aripiprazole 5 mg a day as first choice, with second choices in no particular order: the dopamine agonists cabergoline, bromocriptine and amantadine; Peony Glycyrrhiza Decoction; and metformin 2.5 to 3 g a day. A footnote qualifies the first choice: aripiprazole may not normalise prolactin in amisulpride-induced hyperprolactinaemia. That metformin figure derives from a study of diabetic women on antipsychotics, a narrow population. In the paediatric literature aripiprazole 5 to 15 mg per day is a prolactin-lowering partial dopamine agonist option in children and adolescents, effective without worsening psychosis.

TRAP **Quoting one aripiprazole dose, or one metformin dose, without its indication.** The 3 mg per day and 6 mg per day of the prose and the 5 mg a day of the summary box are read in clean passages; neither is a rendering artefact, and all sit far below antipsychotic doses. All differ again from the 5 to 15 mg per day of adjunctive aripiprazole for weight gain, and that range is not unique to weight gain either: the paediatric prolactin-lowering dose is 5 to 15 mg per day as well, so the number alone settles nothing and the indication and the age group have to be carried with it. Metformin is worse: 500 to 2000 mg per day for weight gain against 2.5 to 3 g per day for prolactin lowering. Name the indication first and the number second, and flag the discrepancy rather than smoothing it away.

IN INDIA

Indian consultation-liaison guidance takes a lower and simpler line, defining hyperprolactinaemia as a serum prolactin above 20 ng/ml, with no separate ceiling for women. Its ranking diverges at the low end: phenothiazines, butyrophenones, risperidone and amisulpride carry the greatest risk; olanzapine and quetiapine modest risk; ziprasidone and aripiprazole low risk, whereas the Maudsley table classes quetiapine as prolactin-sparing; both agree at the high-risk end. It adds that potency of D2 blockade, increasing age and female sex all raise risk, that antidepressants, opioids, antiemetics and antihypertensives are major contributors, and that the strategies are dose reduction, changing agent, a partial D2 agonist such as aripiprazole, or a dopamine agonist, naming bromocriptine. Which of the dopamine agonists is the more available or the cheaper in Indian practice is not addressed by any source drawn on here, so no preference between bromocriptine and cabergoline is stated.

In both lanes the first move is to change the antipsychotic, not add a drug; and if a drug is added, name the indication before the number: metformin at 500 to 2000 mg per day for weight gain but 2.5 to 3 g per day for prolactin, aripiprazole at 5 to 15 mg per day for weight gain, single-figure doses for prolactin in adults, and 5 to 15 mg per day for prolactin again in children and adolescents, so one range serves two indications and only the indication and the age group tell them apart.

18 · Diabetes, depression and hypoglycaemia

Two questions live inside this topic and they are examined differently. The first is epidemiological and prognostic: the bidirectional relationship of diabetes and depression keeps the two travelling together, and the comorbid patient dies earlier than either illness alone predicts. The second is acute and physiological: hypoglycaemia mimics panic before it is recognisably neurological.

The halves connect through treatment, which is where the psychiatrist's leverage lies.

12%	10 TO 15%	50%	70 mg/dL
MAJOR DEPRESSION, POOLED ACROSS 39 STUDIES	DEPRESSIVE DISORDER IN DIABETES, ROUGHLY TWICE THE AVERAGE POPULATION	INCREASED MORTALITY WITH COMORBID DEPRESSION	GLUCOSE AT WHICH THE FIRST PHYSIOLOGICAL RESPONSE BEGINS

18.1 The bidirectional relationship of diabetes and depression, unequally evidenced

Type 2 diabetes and major depression stand in a **bidirectional relationship** in disease incidence, but the arms are not evidenced to the same standard. Studies consistently report an increased risk of major depression in type 2 diabetes, with diabetes doubling the odds of comorbid depression and raising the likelihood of an antidepressant being prescribed, though estimates vary with population, measures and criteria. The depression-to-diabetes arm, an increased risk of insulin resistance and of type 2 diabetes, is asserted flatly, without an effect size, so the honest viva formulation of the bidirectional relationship of diabetes and depression is asymmetric; a remembered relative risk for it fills a blank the sources leave open. Major depression also

contributes significantly to hospital admission and to death in type 2 diabetes, and selective serotonin reuptake inhibition has been reported to improve insulin sensitivity.

The loop is sustained clinically, not molecularly. Depression has a negative impact on metabolic control, and poor metabolic control may in turn worsen depression: a feedback loop the psychiatrist can enter, because the entry point is the depression.

18.2 How common, and in which population

Prevalence figures mean nothing without their denominators. Reported rates of comorbid depressive symptoms in diabetic patients range from 9 to 60%, a spread the source attributes explicitly to differences in study design and screening method rather than real variation between populations. A mid-point would manufacture a number the literature declines to give. Three better-specified estimates:

SOURCE AND POPULATION	FIGURE	WHAT IT IS BOUND TO
Multiple studies, people with diabetes (predominantly Western, higher-income samples)	10 to 15% prevalence of depressive disorders, roughly twice that of the average population	Band and doubling travel together; the band alone loses the claim.
Meta-analysis pooling 39 studies of patients with diabetes	**12%** prevalence of major depression	A twofold increased prevalence against controls was shown in the controlled subset only.
United States health maintenance organisation cohort, approximately 4,800 adult patients with predominantly type 2 diabetes (96%)	12% met criteria for major depression and 8.5% for minor depression	A single insured cohort; the profile below travels no further.

Depression here also shades below the diagnostic threshold: perhaps as many as 25% of people with diabetes experience **subsyndromal depressive symptoms**. They are missed by categorical instruments and are the group in whom self-management quietly deteriorates.

The same cohort supplies a phenotype. Compared with diabetes alone, those with both conditions were younger, more often female, less educated, carried significantly higher medical comorbidity and more diabetes complications, had a longer duration of diabetes, higher glycosylated haemoglobin and body mass index values, and were more often smokers. They had also developed diabetes an average of 5 years earlier. Note what **minor depression** is doing: at 8.5% it is not a rounding error, and it reappears below as an independent contributor to mortality.

IN INDIA

Nearly all the data above are Western and higher-income, and the little evidence from elsewhere suggests the burden here is not smaller. A review from South Asia found the prevalence of type 2 diabetes with elevated depressive symptoms ranged from 11.6% to 67.5%; a meta-analysis from Iran estimated depression in 54% of patients with diabetes; a Saudi study of children with type 1 diabetes reported 27%. Indian investigators found a similarly high rate, with depression associated with younger age, female sex, low socioeconomic status and poor self-management. Two cautions. The standard reference states the Indian finding qualitatively and attaches no figure, so there is no Indian prevalence number to quote. And that risk profile is the patient least likely to return in a fee-for-service setting: screen at the diabetes contact, not on referral.

18.3 Not an episode but a course

Depression here tends to be chronic or recurrent, not self-limiting. In the same insured cohort, 63% of patients with diabetes and major depression had two years or longer of significant depressive symptoms. A separate study, and this is a different population that must not be merged with the first, found that 79% of patients with diabetes and major depression relapsed over a five-year follow-up, with a mean of four episodes per patient.

The first says most comorbid depression met at a clinic visit is a long-standing state, not a fresh episode, which reframes the patient described as unmotivated. The second says remission is not the endpoint. Both point at one decision: maintenance and structured review, not an acute intervention discharged on response.

18.4 What the comorbidity costs

Comorbid depression tracks hard endpoints. A meta-analysis of 22 studies involving more than a million participants found depression associated with increased risk of both macrovascular complications, hazard ratio 1.38 (95% CI 1.30 to 1.47), and microvascular complications, hazard ratio 1.33 (95% CI 1.25 to 1.41), and found the relationship bidirectional at the complication level, in that diabetes complications themselves increased the risk of depression. This is the better-quantified statement of the bidirectional relationship of diabetes and depression, with effect sizes on both arms.

Mortality is where the numbers get garbled: four comparisons in one passage, each with its own denominator.

- A systematic review of 10 longitudinal studies: comorbid depression carried a **50% increased mortality risk**.
- An older-adult study: depression or diabetes each raised mortality by almost one-third, while respondents with both had a fourfold increase compared with those who had neither.
- The Pathways Epidemiology Study, adjusted for age, gender and socioeconomic factors: comorbid major depression carried a twofold increase in mortality over three years against non-depressed patients; minor depression raised mortality risk by 67%.

- The same cohort, further adjusted for diabetes and medical disease severity and for adverse health behaviours including smoking, obesity, sedentary lifestyle and glycosylated haemoglobin: major and minor depression remained associated with an approximately 50% increase in mortality.

The fourth line is the one to carry. If depression killed only through smoking, inactivity and neglected glucose, adjusting for those behaviours should abolish the excess. Roughly half survives, which is why comorbid depression is an independent prognostic factor rather than a marker of poor self-care. The prescribing literature converges on the same magnitude, a 50% increased risk of mortality alongside a high burden of cardiovascular risk factors.

18.5 Shared mechanisms, held loosely

The closeness of the association invites a common-cause explanation, and several are on offer. The commonest is psychological: chronic disease burden predisposes to depression, or depression erodes motivation and worsens self-management. Shared biology is proposed alongside it. Hyperglycaemia can affect mood states, and high plasma glucose in the brain can increase oxidative stress, producing neuronal damage and, on this account, depression. The microvascular changes of diabetes are identifiable in the brain and may predispose to depression independently. Type 2 diabetes is associated with **low-grade systemic inflammation**, proposed as the substrate of the depression seen in diabetes, and others have suggested calcium and cyclic AMP signalling as the cellular mechanism.

The source supplies the essential qualifier itself: most of this work is cross-sectional and limited to correlation, and longitudinal research is needed before any of it can be called established. Write them as candidate mechanisms; an answer asserting inflammation as the mechanism has overstated the evidence.

18.6 Making the diagnosis of diabetes: two thresholds that do not agree

The screening decision often falls to the psychiatrist. Three clinical measures may be used: a fasting plasma glucose of 126 mg/dL or above; a non-fasting plasma glucose of 200 mg/dL or above in a person with symptoms of diabetes; or an abnormal oral glucose tolerance test with a two-hour value of 200 mg/dL or above. Alongside these, **glycosylated haemoglobin** has become a standard method both of diagnosis and of monitoring over time, reflecting glycaemic control over the preceding 2 to 3 months.

The guideline written for patients on antipsychotic medication states the same rules with different edges: a fasting blood glucose higher than 125 mg/dL, with fasting defined as no caloric intake for at least 8 hours, or alternatively a glycosylated haemoglobin of 6.5% or greater; an oral glucose tolerance test or a random blood glucose of at least 200 mg/dL in conjunction with a hyperglycaemic crisis or classic hyperglycaemic symptoms are also acceptable. Crucially, it adds that results should be confirmed by repeat testing unless unequivocal hyperglycaemia is present.

TRAP **There is no single harmonised glycosylated haemoglobin rule, and a question written around a value exactly on the threshold is testing whether you know that.** The comprehensive textbook states that two values over 6.5% are diagnostic: two measurements, strict inequality. The practice guideline, citing the American Diabetes Association, states that a single value of 6.5% or greater can be used: one measurement, inclusive threshold, with confirmation by repeat testing unless hyperglycaemia is unequivocal. A reading exactly on the threshold satisfies the second and fails the first. The fasting criterion repeats the pattern, 126 mg/dL or above against higher than 125 mg/dL: identical for whole numbers, differently expressed. Do not resolve this by picking a value; the point is that confirmation is required.

One further rule marks the limit of the convenient test. In haemoglobinopathies, or conditions associated with increased red blood cell turnover, fasting blood glucose should be used rather than glycosylated haemoglobin; the conditions named are second-trimester or third-trimester pregnancy, haemodialysis, recent blood loss, transfusion and erythropoietin therapy. These criteria belong with the account of metabolic syndrome, obesity and mortality in serious mental illness; the screening schedule, with the metabolic monitoring protocols.

18.7 Disordered eating and insulin omission in type 1 diabetes

This small topic punches above its weight, and belongs to type 1 diabetes. A meta-analysis reports eating disorders in 8.2% of patients with type 1 diabetes compared with 2.8% of controls. What makes the population distinctive is **insulin omission**, a purging method available only to a patient who injects insulin, found in **up to 40% of type 1 diabetic patients**. Insulin is withheld to lose weight; risk factors are female sex and higher body mass index. Consider an eating disorder in any patient with poor glycaemic control or multiple episodes of diabetic ketoacidosis. Neither figure carries across to type 2 diabetes, where the behaviour has no equivalent mechanism.

PITFALL

Recurrent ketoacidosis in a young woman with type 1 diabetes is routinely written up as non-adherence, and the note recommends education. Non-adherence and deliberate insulin omission look identical on a discharge summary and are entirely different problems: one needs teaching, the other an eating disorder assessment. The discriminating questions are about weight, shape and the function the omission serves, and nobody asks unless the possibility is raised.

18.8 Hypoglycaemia: the cascade, the thresholds and the symptom sequence

In the absence of prolonged fasting, healthy people almost never have blood glucose below 3.5 mM. As glucose falls, defence is mounted in stages. The pharmacological account places **reduction of endogenous insulin secretion** first, occurring at a plasma glucose of **about 70 mg/dL (3.9 mM)**; thereafter the **counter-regulatory hormones** are released, named as glucagon, epinephrine, norepinephrine, growth hormone and cortisol. Symptoms typically become manifest at 60 to 70 mg/dL (3.3 to 3.9 mM).

TRAP **Two standard texts order the physiology differently, and an answer built on the wrong ordering is wrong even with the right number.** The pharmacology text puts falling endogenous insulin secretion first and hormone release second, naming five counter-regulatory hormones. The clinical methods text opens with hormone release, stating that physiological responses begin at a plasma glucose of around 3.8 mmol/L with the release of counter-regulatory hormones such as glucagon or adrenaline, and names two. The thresholds are the same, around 3.8 mmol/L against about 3.9 mM; the disagreement is what happens first. State that both accounts exist, and prefer the fuller list when asked to enumerate.

The sequence is ordered by glucose level, not severity of distress: autonomic symptoms come first, generated at a higher glucose than neuroglycopenic ones.

STAGE	FEATURES	WHY THE ORDER MATTERS
Autonomic, at higher glucose	Sweating, hunger, paraesthesiae, palpitations, tremor and anxiety, with agitation and blurred vision described alongside them	The phenomenology of a panic attack, hence the anxiety differential. Most patients recognise these symptoms and treat themselves rapidly.
Neuroglycopenic, at lower glucose	Difficulty concentrating, confusion, weakness, drowsiness, a feeling of warmth, dizziness, blurred vision and loss of consciousness; impaired intellectual activity and diminished psychomotor skills	Cerebral glucose is now low. The patient can no longer reliably self-rescue; the presentation is cognitive.
Severe	Seizure and coma; severe agitation, confusion, coma and epileptiform seizures	Hypoglycaemia can present as an acute behavioural disturbance, not a faint.

One qualification undoes that reassurance. **Loss of awareness of hypoglycaemia** is sometimes a problem in diabetic patients with autonomic neuropathy, and in insulin-treated diabetes leads to unexpected hypoglycaemia: the warning tier is absent, and the patient reaches the neuroglycopenic stage without having felt the autonomic one. Those at risk should be warned against driving motor vehicles and taught the early symptoms of this potentially serious complication. The standard texts give no glucose value at which neuroglycopenic symptoms begin, and none should be invented.

18.9 Hypoglycaemia and psychiatric presentations: the differential, and managing the pair

The link between hypoglycaemia and psychiatric presentations earns its place twice over: as a mimic and as a hazard. As a mimic, it appears in the standard list of medical conditions associated with anxiety, in the metabolic group alongside hyperkalaemia, hyperthermia, hyponatraemia and vitamin deficiencies; hyponatraemia belongs to the account of SIADH and hyponatraemia. The discriminating feature is temporal: phaeochromocytoma and hypoglycaemia usually produce **episodic symptoms** and so are likelier to mimic a phobic or panic disorder, though both belong in the differential of generalised anxiety disorder, with examination and laboratory tests where there is doubt.

MNEMONIC

Treatment, Tumour, Taken

The commonest scenarios producing hypoglycaemia. **Treatment**: treatment of diabetes with insulin or with oral agents that promote insulin secretion. **Tumour**: inappropriate endogenous insulin from a pancreatic islet tumour, or insulin-like growth factors from non-islet tumours such as hepatomas or sarcomas. **Taken**: purposeful or inadvertent use of a glucose-lowering agent by a person without diabetes, the factitious route and the psychiatrically important one. The first and third can occur in either the fasting or the fed state, whereas hypoglycaemia from a neoplasm occurs in the fasting state. Within treated diabetes it may follow an inappropriately large insulin dose, a mismatch between peak insulin delivery and food intake, or added factors that increase insulin sensitivity, such as adrenal or pituitary insufficiency, or that increase insulin-independent glucose uptake, such as exercise.

As a hazard, it complicates glucose-lowering treatment, particularly insulin, and leaves a cognitive trace: severe hypoglycaemic episodes, meaning those associated with emergency department attendance or hospitalisation, have increased the risk of dementia in patients with diabetes. The association is bound to episodes of that severity and carries no effect size. The dementia syndrome and its workup belong to the Dementias and neurocognitive disorders notes.

RULE **Every patient presenting with coma or a reduced conscious level must have their blood glucose checked immediately.** Patients with diabetes can present with acute metabolic decompensation, leading to diabetic ketoacidosis, hyperosmolar hyperglycaemic syndrome or, in treated patients, hypoglycaemia. There is no exemption for a patient who already holds a psychiatric diagnosis. The wider assessment of the acutely confused patient belongs to the Stroke, TBI, delirium and focal brain syndromes notes; the capillary glucose comes first, and is not delegated.

LOCUS. Mostly outpatient and shared with a physician or diabetologist, the psychiatrist contributing recognition, sequencing and monitoring rather than glycaemic prescribing. The emergency locus is the rule above; inpatient care belongs to the patient whose depression has made self-management unsafe, or whose recurrent ketoacidosis needs both medical stabilisation and an eating disorder assessment. The authorities are The Maudsley Prescribing Guidelines for antidepressant selection, the American Psychiatric Association practice guideline for screening thresholds, and Kaplan and Sadock's Comprehensive Textbook of Psychiatry for epidemiology, prognosis and the eating disorder overlap.

FOCUS changes with phase. The acute target is glucose itself and the safety of a patient confused or agitated from it. The short-term target is the depressive episode on its own merits, with glycaemic control watched as it moves. The mid-term target is the recurrence risk above, making structured review the default. The long-term targets are the complication and mortality trajectories and the cognitive cost of severe hypoglycaemia.

MODUS. Every patient diagnosed with depression should be screened for diabetes. In those who are diabetic, selective serotonin reuptake inhibitors are used first line, with data supporting

sertraline, escitalopram and fluoxetine; serotonin and noradrenaline reuptake inhibitors are also likely to be safe but have fewer supporting data; agomelatine seems promising on limited data; and tricyclic antidepressants and monoamine oxidase inhibitors should be avoided where possible because of their effects on weight and glucose homeostasis. That last caution has a sharp edge here, because **irreversible monoamine oxidase inhibitors** have a tendency to cause extreme hypoglycaemic episodes as well as weight gain, an effect explicitly not attributed to moclobemide. Psychological and social modes carry the self-management burden the epidemiology keeps pointing at; the pharmacology of the antidepressants belongs to the Mood disorders: pharmacotherapy notes.

GOOD TO REMEMBER

The monitoring instruction is more specific than general vigilance. Blood glucose and glycosylated haemoglobin should be monitored carefully when antidepressant treatment is started, when the dose is changed, and after it is discontinued. Discontinuation is the forgotten one: the glucose effect outlasts the prescription.

Depression in diabetes is common, chronic, recurrent and independently lethal, carrying an approximately 50% increase in mortality that survives adjustment for health behaviours; and in any patient with reduced consciousness, seizure or acute behavioural disturbance, capillary glucose is measured immediately and before the presentation is called psychiatric.

19 · Metabolic syndrome, obesity and mortality in serious mental illness

This is where psychiatry stops being a specialty of the mind alone. Patients with schizophrenia and related psychoses do not, in the main, die of psychosis. They die of the cardiovascular and metabolic diseases that kill everyone else, only earlier and after less vigorous medical attention. Metabolic syndrome names the cluster that antipsychotic treatment aggravates, and it is the physiology of that premature death.

34%	40 TO OVER 50%	TWICE	10 TO 25 YEARS
METABOLIC SYNDROME, US ADULTS, GENERAL POPULATION	METABOLIC SYNDROME IN SCHIZOPHRENIA	OBESITY PREVALENCE IN SCHIZOPHRENIA VERSUS GENERAL POPULATION	LIFE-EXPECTANCY REDUCTION IN SERIOUS MENTAL DISORDER

19.1 The construct, and the bodies that own the definition

Metabolic syndrome, or syndrome X, was delineated by **Gerald Reaven** and remains contested in its detail. It is not a disease in the way diabetes is but a way of noticing that these abnormalities travel together and jointly predict outcomes no single component predicts as well alone.

Companion to Psychiatric Studies counts four major features: increased abdominal or visceral adiposity seen as raised waist circumference, atherogenic dyslipidaemia, hypertension, and impaired fasting glucose or overt diabetes. Four and not five because **atherogenic dyslipidaemia** bundles low high-density lipoprotein and raised triglycerides into one: presentation, not content.

Kaplan and Sadock trace the most widely accepted definition to the third adult treatment panel (ATP III) of the National Cholesterol Education Program, then record two modifications: the adapted panel (ATP III-A) proposed by the American Heart Association decreased the cutoff for fasting glucose, and the International Diabetes Federation stressed waist circumference by modifying the parameters to a lower cutoff. Companion agrees: the American Heart Association set requires a slightly lower fasting glucose, and the International Diabetic Association criteria, the most pragmatic if the least rigorous, are based solely on waist circumference.

19.2 The operational criteria, and the two structures that disagree

The citable set is the harmonised definition in the American Psychiatric Association practice guideline, attributed to Alberti et al. 2009: **at least three of five** risk factors, none mandatory. Two features are load-bearing and routinely dropped. First, the waist threshold is the source's own declared United States and Canada value, not a universal one. Second, each non-waist criterion is met by the number or by drug treatment for that abnormality: a fibrate, nicotinic acid or high-dose omega-3 fatty acids for the lipid criteria, an antihypertensive in a patient with a history of hypertension, or a glucose-lowering drug, so a patient controlled on medication can satisfy three criteria with every printed laboratory value normal.

A structurally different set circulates. Appendix 2 of the Maudsley Guidelines on Advanced Prescribing in Psychosis makes waist obligatory and requires only two of the remaining four, naming no issuing body.

FEATURE	HARMONISED SET (AMERICAN PSYCHIATRIC ASSOCIATION PRACTICE GUIDELINE)	WAIST-MANDATORY SET (MAUDSLEY GUIDELINES ON ADVANCED PRESCRIBING IN PSYCHOSIS, APPENDIX 2)
Structure	Any three of five, nothing mandatory	Waist obligatory plus any two of four
Waist, men	Above 102 cm (40.2 inches), a declared United States and Canada value	94 cm (37 inches) or more
Waist, women	Above 88 cm (34.6 inches)	80 cm (31.5 inches) or more
Triglycerides	150 mg/dL or more	1.7 mmol/L or more
High-density lipoprotein	Below 40 mg/dL men, below 50 mg/dL women	Below 1.03 mmol/L men, below 1.29 mmol/L women
Blood pressure	Systolic 130 mmHg or more, diastolic 85 mmHg or more, or both	130/85 mmHg or more
Fasting glucose	100 mg/dL or more	Above 5.5 mmol/L

MNEMONIC

Wide, Thick, Low, High, Sweet

In the harmonised set's own order: **Wide** waist, **Thick** blood (triglycerides), **Low** high-density lipoprotein, **High** blood pressure, **Sweet** blood (fasting glucose). Any three make the syndrome, and for the last four a drug treating the abnormality counts as the abnormality.

TRAP **Treating "the metabolic syndrome criteria" as one thing.** The disagreement is structural before it is numerical. Raised triglycerides, low high-density lipoprotein and raised blood pressure with a normal waist *is* the syndrome by the any three of five rule and *is not* by the waist-obligatory rule. Name the criteria set before applying it.

The units do not convert: the harmonised threshold for elevated fasting glucose, 100 mg/dL or more, appears in the Maudsley appendix as a strict inequality in different units, and converting between them would be the chapter's arithmetic, not a source's threshold. It is also the threshold that *moved*: both sources naming the American Heart Association adaptation as the one that lowered it give no value it was lowered from, so the original ATP III cutoff is not recoverable here. Tasman Psychiatry states the same rule in imperial units, abdominal obesity as a waist of 35 inches or more for women or 40 inches or more for men.

19.3 The definition you apply decides the prevalence you find

Kaplan and Sadock put general-population prevalence at approximately 34% of United States adults, similar to other developed countries, up from 25% in 1988, concentrated in non-Hispanic black females, and also associated with lower education and older age. Most relevant studies put prevalence in schizophrenia at about 40% to over 50%, with no criteria set named for that range. That range then circulates as the prevalence of metabolic syndrome in severe mental illness, a label wider than the schizophrenia samples it was measured in; the widening is the reader's and not the source's.

The CATIE baseline sample shows why. In the same patients, National Cholesterol Education Program criteria gave 40.9% and American Heart Association criteria gave 42.7% (McEvoy et al 2005). Change the instrument again and the spread widens: using waist circumference alone, prevalence in females was double that in males, 73.4% against 36.6%, and even accounting for body mass index, males were 85% more likely to have metabolic syndrome while the comparable figure in females was 137%.

RULE **A prevalence of metabolic syndrome quoted without its criteria set and its population is uninterpretable.** Two criteria sets applied to the same patients differ by under two percentage points, 40.9% against 42.7%. The wide spread comes from elsewhere: using waist circumference alone, that same cohort runs 36.6% in males against 73.4% in females, a split by sex within one method rather than a gap between methods. Quote neither figure without saying which criteria set was applied and to whom.

19.4 What the syndrome actually buys in risk

Prospective studies show that metabolic syndrome raises the risk of developing coronary vascular disease by over 50% and the risk of type 2 diabetes by approximately threefold compared with individuals without metabolic syndrome. Both figures carry that comparator inside them. The **approximately threefold** diabetes risk is why the cluster is called prediabetic rather than a risk marker, and the thread is **insulin resistance**, which is why these abnormalities cluster at all. In the United States over 96 million adults, or 1 in 3 adults aged 18 or older, have prediabetes.

IN INDIA

Every prevalence figure here is United States or Western, the sharpest limitation in this topic for Indian practice. The harmonised waist threshold is its own source's United States and Canada value, and no South Asian threshold appears in these references at all; since central adiposity at a given waist means different things across populations, a North American cut-point applied to Indian patients is not neutral. There is likewise no Indian general-population prevalence, no Indian prevalence in serious mental illness, and no Indian or South Asian mortality-gap figure. Record the waist in centimetres rather than only the verdict against an imported threshold, and the number stays usable when a local cut-point arrives. Lipids and glucose cost the patient money; the tape, the scale and the cuff cost nothing, and the free measurements should never be the omitted ones.

19.5 Obesity: the thresholds, and the number that triggers action

Obesity is the component the psychiatrist sees first. It is about twice as prevalent in people with schizophrenia as in the general population, and substantial evidence attributes part of that excess to psychotropic treatment, particularly certain antipsychotics. The drug-by-drug ranking belongs with antipsychotic metabolic effects, and drug selection with the Schizophrenia: management, antipsychotics and adverse effects notes.

Body mass index is weight in kilograms divided by the square of height in metres. Kaplan and Sadock give greater than 25 kg/m2 for overweight and greater than 30 kg/m2 for obesity, and record approximately 42% of United States adults obese as of 2019, up from 30% in 2000, with severe obesity rising from about 5% to 9%: Centers for Disease Control figures. The American Psychiatric Association practice guideline bands the same construct as 25 to 29.9 for overweight and 30 or higher for obesity, and adds waist as a risk indicator at above 35 inches for women and above 40 inches for men.

In patients on antipsychotic medication, except for those with a body mass index below 18.5, an increase in body mass index of 1 unit suggests a need for intervention: closer weight monitoring, a weight management programme, adjunctive treatment to reduce weight, or changing the antipsychotic.

GOOD TO REMEMBER

The threshold for action is **an increase in body mass index of 1 unit**, not a crossing of the obesity line. It is set that low because weight gained early on an antipsychotic is rarely lost later, so waiting for formal obesity wastes the whole period in which the trajectory could have been altered.

Why body mass index rather than the waist that the metabolic syndrome criteria make central? Stahl's The Clozapine Handbook is explicit: waist circumference shows significant variability between examiners to an extent that does not exist with body mass index, so changes in body mass index serve as a more accurate measure. Diagnosis asks whether a threshold was crossed once, monitoring asks whether a number is moving, and inter-rater noise wrecks the second question while barely touching the first.

19.6 Diabetes and cardiovascular comorbidity in serious mental illness

Kaplan and Sadock count almost 100 systematic reviews and meta-analyses of physical health comorbidity in mental illness since 2000, of which 15 concerned schizophrenia between 2000 and 2019, spanning cardiovascular disease, cancer, diabetes, metabolic syndrome, osteoporosis and arthritis, and conclude that physical comorbidities are the rule and not the exception.

For type 2 diabetes, prevalence is substantially higher than in the general population, with estimates ranging from two to four times higher. Three pooled estimates in the same passage must not be collapsed: Stubbs and colleagues reviewed 25 studies and found 9.5%; Vancampfort and colleagues pooled 22 studies for 11.5%, with a risk of nearly double the general population; and Ward and Druss estimated a median prevalence of 13%. The source attributes the spread to study design, detection and sample size, and leaves standing the tension between the two-to-four-times headline and a pooled estimate of nearly double. The diabetes-depression relationship and hypoglycaemia are taught with diabetes, depression and hypoglycaemia.

Cardiovascular disease is the leading cause of mortality in schizophrenia. Fan and colleagues reviewed 13 studies with 422,698 patients with schizophrenia carrying cardiovascular outcome data and calculated pooled relative risks:

- 1.53 for cardiovascular disease overall
- 1.2 for coronary heart disease
- 1.7 for stroke
- 1.81 for congestive heart failure

Correll and colleagues, in more than 3 million patients, give prevalence rather than risk: a pooled cardiovascular disease prevalence of 9.9% across all patients with serious mental illness, significantly higher than normal controls. The populations differ, Fan being schizophrenia-specific and Correll covering serious mental illness as a whole. Correll's effect estimate is corrupt in the available text, its point estimate falling outside its own confidence interval, and is not reproduced.

Bipolar disorder enters these data at one point only, and it is cardiovascular: patients with schizophrenia and bipolar disorder are twice as likely to die from cardiovascular disease as the general population. The same passage makes it a health-systems problem, since patients with diabetes and schizophrenia are treated less aggressively for their cardiovascular risk than patients with diabetes and no mental disorder, and carry an elevated rate of depression. Part of the gap is not a property of the illness but of the care it attracts.

19.7 The mortality gap: how many years, and for whom

DSM-5-TR asserts the reduction and declines to quantify it, attributing it to associated medical conditions, with weight gain, diabetes, metabolic syndrome, cardiovascular and pulmonary disease all commoner than in the general population, and adds a hedge: a shared vulnerability for psychosis and medical conditions may explain some of the medical comorbidity.

SOURCE	POPULATION AS STATED	ESTIMATE
Kaplan and Sadock's Comprehensive Textbook of Psychiatry	Schizophrenia *and other chronic psychotic disorders*	About 10 to 15 years prematurely; men about 15, women on average 12
Shorter Oxford Textbook of Psychiatry	Schizophrenia	Men approximately 20 years, women 15 years
Stahl's Essential Psychopharmacology	Schizophrenia	"May be" 20 to 30 years shorter
Child and Adolescent Psychopathology	Serious mental disorder, undifferentiated	Averaging 10 to 25 years (Chang et al., 2011)
DSM-5-TR	Schizophrenia	Reduction asserted, no figure given

TRAP **"People with serious mental illness die twenty years early."** No single source supports that sentence. The two sex-disaggregated estimates contradict each other for the same illness, and the only estimate whose population genuinely *is* serious mental illness is the broadest and vaguest, at **10 to 25 years**. Give the range, name the sources, and say why it varies: different populations, different eras, different metrics.

Excess mortality was originally reported by Odegaard in Norway in 1951, and the estimate has moved: an early meta-analysis by Harris and Barraclough (1998) found all-cause mortality increased 1.6-fold, while subsequent studies found a higher figure, typically around threefold, and suggested that the mortality gap is increasing (Saha et al., 2007). The **mortality gap** also widens unevenly: all-cause mortality is increased several-fold further in patients who have comorbid alcohol and substance misuse (Hjorthoj et al., 2015). Physical illness rarely comes singly: perhaps as many as two-thirds of these patients with comorbidity have more than one physical disease, that is, multimorbidity, and they are less likely to receive and benefit from appropriate medical treatment and more likely to experience unfavourable clinical outcomes.

19.8 What patients actually die of, and the contradiction the sources leave standing

The references split into two irreconcilable camps. Kaplan and Sadock, reporting the Global Burden of Disease analysis, put the weight on physical disease: cardiovascular disease alone accounted for more than 33% of premature deaths, while suicide, homicide and accidents together accounted for less than 15% of the excess deaths, with Brown and colleagues, following 370 schizophrenia patients over 25 years, finding 81% died from medical diseases. The Shorter Oxford Textbook of Psychiatry states the opposite: about 60% of the excess early mortality in schizophrenia is accounted for by unnatural causes, meaning suicide and accidents, and the rest by natural causes, chiefly cardiovascular, attributed to smoking, poor diet, sedentary lifestyle,

obesity and type 2 diabetes, with lower rates of diagnosis and less active treatment also contributing (Shiers et al., 2015).

Under 15% against about 60% for what looks like the same quantity, and neither passage can be dismissed. Two reconciling possibilities, neither asserted: a share of *all* premature deaths is not a share of *excess* deaths over general-population expectation; and cohorts differ in era and in weighting towards early illness, when suicide risk is concentrated. A merged sentence would create a fact neither source supports.

The **standardised mortality ratios** are less contested. Cohort studies consistently show all-cause mortality two to four times higher than the general population. In Osby and colleagues' sample of 7,784 inpatients, the ratio for death from any cause was 2.8 in men and 2.4 in women, an inpatient sample, which the source notes tends to yield smaller increases. In Olfson and colleagues' study of over 1 million United States Medicaid patients, those with schizophrenia were 3.7 times more likely to die during follow-up, with cardiovascular disease at a ratio of 3.6 and the highest mortality rate of any single cause, 403.2 per 100,000 person-years, and elevated ratios for chronic obstructive pulmonary disease at 9.9, influenza and pneumonia at 7.0, and diabetes mellitus at 4.2 with a rate of 61.8 per 100,000 person-years. Against most clinicians' intuition, accidental deaths at 119.7 per 100,000 person-years accounted for more than twice the deaths attributable to suicide at 52.0 per 100,000 person-years.

Suicide nonetheless contributes heavily: between 25% and 50% of schizophrenia patients attempt suicide and up to 10% eventually succeed, contributing to a mortality rate eight times greater than that of the general population. That eight-fold figure belongs to the suicide sentence and is not a free-standing all-cause ratio.

19.9 What follows for practice

The gap is partly iatrogenic, partly behavioural and partly a failure of medical care, and only the last is quickly fixable. LOCUS. Almost all of this work happens in the outpatient psychiatric clinic, where the psychiatric contact is often the patient's only reliable medical contact of the year. Admission is the opportunity for the baseline anthropometry and bloods that outpatient attendance never quite delivers; emergency presentation belongs to the acute diabetic and cardiac decompensations.

FOCUS. Acutely the target is the measurement itself: nothing can be acted on that has not been recorded. In the short term it is trajectory rather than category; in the medium term the individual components, each separately treatable; in the long term, the mortality gap.

MODUS. The American Psychiatric Association practice guideline supplies the metabolic syndrome criteria used here and the body mass index action threshold; the Maudsley Guidelines on Advanced Prescribing in Psychosis supply the alternative waist-obligatory set used in psychosis services. Management of established disturbance is set out with the management of antipsychotic metabolic disturbance, and the schedule with the metabolic monitoring protocols. Diet, activity and above all smoking cessation address the causes the epidemiology identifies.

PITFALL

Concluding that because antipsychotics cause metabolic syndrome, less antipsychotic means longer life. The data run the other way: moderate levels of antipsychotic use are associated with the lowest mortality in schizophrenia, and even high doses are associated with lower mortality than no treatment at all (Torniainen et al., 2015). This is observational association rather than demonstrated causation and no dose is specified, but the direction is not what the metabolic story predicts: untreated psychosis carries its own lethality through suicide, accident and disengagement from medical care. From the same passage, the incidence of, and mortality from, some cancers and autoimmune disorders is lower than expected in individuals with schizophrenia and in their relatives (Catts et al., 2008), by unknown mechanisms.

Metabolic syndrome is any three of five risk factors in one definition and an obligatory waist plus two others in another, so no prevalence figure means anything without its criteria set; and every standard source agrees life expectancy in schizophrenia is substantially reduced but not on the years or the causes, so the defensible answer gives the range with its populations, never a single confident number.

20 · SIADH and hyponatraemia: mechanisms, psychiatric drug causes and management

Hyponatraemia is the electrolyte disturbance this specialty actually causes. Nearly every class we prescribe has been implicated, the presentation is a confusional state that looks like the illness being treated, and the drug responsible is usually one started weeks ago. Marks cluster on the definition, the single test separating the two routes to a low sodium, drug-class risk where candidates guess, and a short set of action thresholds. Four sources carry the territory, and where two disagree the disagreement is the question: the Maudsley Prescribing Guidelines for epidemiology, risk and thresholds, Kaplan and Sadock's Comprehensive Textbook of Psychiatry for mechanism and anticonvulsants, Brazis's Localization in Clinical Neurology for the criteria, Goodman and Gilman's The Pharmacological Basis of Therapeutics for treatment.

AS HIGH AS 11%	1 IN 300	30 DAYS	0.004% TO 26.1%
SIADH PREVALENCE, ACUTELY ILL PSYCHIATRIC PATIENTS	HYPONATRAEMIA ADMISSION ON ANY ANTIDEPRESSANT, OLDER ADULTS	USUAL ONSET WINDOW AFTER STARTING	REPORTED ANTIPSYCHOTIC PREVALENCE RANGE

20.1 The defining biochemistry, and where normal ends

The **syndrome of inappropriate secretion of antidiuretic hormone** is defined biochemically, not clinically. Brazis states the conditions, satisfied together, as serum hyposmolarity below 280 mOsm/kg with hyponatraemia below 135 mEq/L, normal renal excretion of sodium, and an inappropriately high urine osmolality without body fluid depletion, with renal and adrenal function normal and serum ADH elevated. His subtitle, elevated ADH release with normal thirst, states what makes the secretion inappropriate: the hormone appears when the signals that normally drive it are absent.

Two elements do the diagnostic work and are the most often dropped. Absence of fluid depletion matters because a dry patient has every reason to secrete the hormone, so this is a euvolaemic diagnosis. Normal renal and adrenal function excludes renal salt wasting and adrenal insufficiency, a hyponatraemia that will not answer to anything done about the psychotropic. The Maudsley gives the normal serum sodium range as **136-145 mmol/L**.

TRAP **The lower limit of normal is 135 or 136 depending on which book you revised from.** Brazis treats a sodium below 135 mEq/L as the abnormality, leaving that value itself normal; the Maudsley range begins at 136 mmol/L, putting the Brazis value already below it. Neither is a misprint, so do not harmonise them and do not assume the paper you are sitting has. The reassurance is clinical: no decision turns on that unit, because every action threshold sits far below it.

20.2 The mechanism, and the one discrimination made at the bedside

Kaplan and Sadock are mechanistically direct: excessive secretion of **arginine vasopressin** increases retention of fluid in the body, and the retained water produces the hyponatraemia. Total body sodium is not the problem; the solvent is in excess. It may develop after injury to the brain or from medication, including phenothiazines, butyrophenones, carbamazepine, oxcarbazepine and SSRIs among many other agents, and the hyponatraemia may produce delirium. Delirium as a syndrome belongs to the Stroke, TBI, delirium and focal brain syndromes notes; here, new confusion in a recently medicated patient calls for electrolytes, not a dose increase.

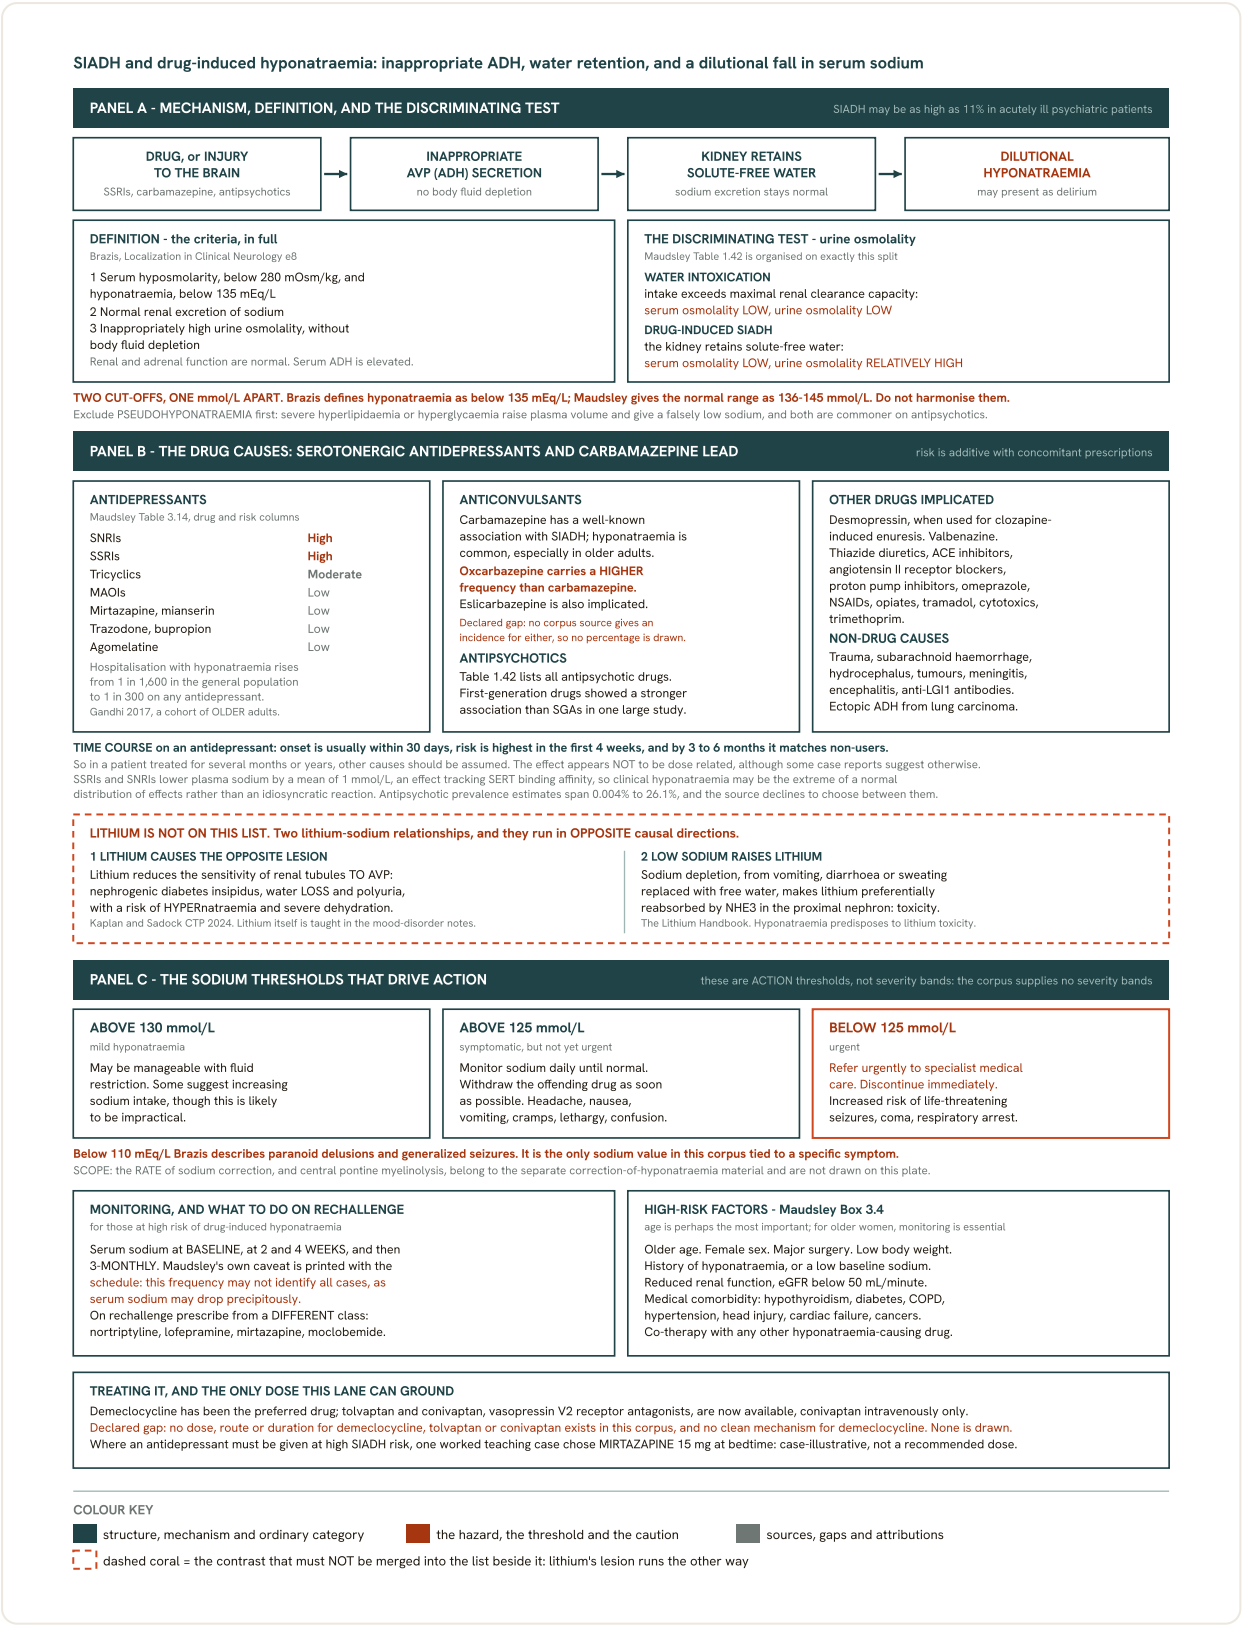


Fig 33.8 SIADH and drug-induced hyponatraemia: the mechanism, the discriminating test, the drug causes, and the sodium thresholds that drive action. Excess AVP makes the kidney retain solute-free water, so sodium falls by dilution while renal sodium excretion stays normal; the bedside discriminator is urine osmolality, low in water intoxication and relatively high in SIADH. Serotonergic antidepressants and carbamazepine lead, and oxcarbazepine carries the higher frequency of the two, though no corpus source attaches an incidence to either. Lithium is deliberately absent: its lesion runs the opposite way. (Maudsley 15th edition; Kaplan and Sadock's CTP 2024; Brazis, Localization in Clinical Neurology e8; Goodman and Gilman; The Lithium Handbook; Geriatric Psychiatry Study Guide.)

Psychiatric patients reach a low sodium by two routes needing opposite logic. In **water intoxication** intake outruns maximal renal clearance and both serum and **urine osmolality** are low. In drug-induced SIADH the kidney retains solute-free water, so serum osmolality is low while the urine stays concentrated: suspect it when a patient is hyponatraemic and the urine is concentrated, that is, when urine osmolality or specific gravity is high. Compulsive drinking in its own right is taken up where this chapter deals with psychogenic polydipsia, the sodium-correction risk and mineral disturbance.

AXIS	WATER INTOXICATION	DRUG-INDUCED SIADH	LITHIUM-ASSOCIATED NEPHROGENIC DIABETES INSIPIDUS
Primary lesion	Intake exceeds renal clearance	Kidney retains solute-free water	Tubules less sensitive to vasopressin
Water moves	In, in excess	Retained under vaso-pressin drive	Out, as polyuria
Serum osmolality	Low	Low	Not characterised in these sources
Urine	Osmolality low	Osmolality relatively high	Cannot be concentrated
Sodium moves	Down	Down	Up, with dehydration risk

RULE **Serum osmolality is low in both psychiatric mechanisms; only the urine tells you which.** A dilute urine says the kidney is working against an intake it cannot match: stop the drinking. A concentrated urine says it is obeying a signal it should not be receiving: find the drug. One paired serum and urine osmolality separates two managements sharing no steps.

20.3 Antidepressants: the best-quantified cause

The Maudsley's SIADH prevalence of as high as 11% in acutely ill psychiatric patients is an upper bound, not a point estimate. As risk, hospitalisation with hyponatraemia rises from 1 in 1,600 in the general population to 1 in 300 on any antidepressant, from a cohort of older adults; quote the population with the number, since the effect is age-loaded and reads as alarming in a young outpatient otherwise.

Onset is usually within 30 days of starting, risk is highest in the first 4 weeks, and falls until by 3 to 6 months it equals that of patients not taking antidepressants. Against intuition, the effect appears not to be dose related, although some case reports suggest otherwise, so dose reduction is not rational and the choice lies between stopping and monitoring. Across the treated population the magnitude is small: a 2022 Japanese study found SSRIs and SNRIs caused an average plasma sodium fall of **1 mmol/L**, absent from the other antidepressants examined and linked to binding affinity for the serotonin transporter. If the whole distribution shifts, clinical hyponatraemia is plausibly its tail rather than an idiosyncratic reaction in a few.

ANTIDEPRESSANT CLASS	LEVEL OF HYPONATRAEMIA RISK
SNRIs	High
SSRIs	High
Tricyclics	Moderate
MAOIs	Low
NaSSAs (mirtazapine, mianserin)	Low
Trazodone	Low
Bupropion	Low
Agomelatine	Low

The narrative agrees: a 2024 meta-analysis suggests risk may be highest with SNRIs, then SSRIs, and lowest with mirtazapine and trazodone, fitting the transporter-affinity finding. The pharmacovigilance literature agrees with neither. One large regulatory database study found the strongest association with mirtazapine, a French study the greatest risk with duloxetine, another a signal for agomelatine, not previously reported to cause hyponatraemia at all.

TRAP **Quoting a spontaneous-reporting signal as though it were a risk estimate.** The Maudsley explains its own divergent findings: disproportionate reporting arises for drugs whose adverse effect is believed rare, confounding by indication cannot be adjusted for because drugs thought low risk go preferentially to patients already at high risk, and concomitant prescriptions distort the signal. The mirtazapine finding reads as an artefact of who receives mirtazapine. Give the table's ranking, then name the artefact. Note too that no antidepressant has been definitively shown not to be linked with hyponatraemia.

GOOD TO REMEMBER

A patient stable on the same antidepressant for a year who becomes hyponatraemic probably has another cause. Risk has fallen to a non-user's by 3 to 6 months, so look for the new diuretic, the chest lesion or the intercurrent illness before stopping a drug that works.

20.4 Antipsychotics, and the anticonvulsant direction candidates reverse

The antipsychotic evidence is broad and far less precise. Case reports and series implicate various phenothiazines, haloperidol, pimozide, risperidone, paliperidone, quetiapine, olanzapine, aripiprazole, cariprazine and clozapine, and one large Swedish study found a stronger association for first-generation than second-generation drugs. The Maudsley tabulates the implicated agents as all antipsychotic drugs. Two further points are examinable: antipsychotic treatment was suggested to be five times more likely than water intoxication to be the cause, and reported prevalence spans 0.004% to 26.1%, the source assuming the truth lies between two widely different extremes. Offer no single figure; the refusal to pick one is the point. Age, sex, comorbidity and polypharmacy are said to matter less here than for antidepressants.

Among anticonvulsant mood stabilisers carbamazepine is the archetype, its association called well known by the Maudsley. Kaplan and Sadock add that hyponatraemia is common with

carbamazepine, particularly in older adults, that baseline electrolytes may reasonably be taken in vulnerable populations, and that neurocognitive change on it should prompt a sodium check, an effect shared with oxcarbazepine and eslicarbazepine.

TRAP **Naming carbamazepine as the worst anticonvulsant offender.** Kaplan and Sadock state that oxcarbazepine is structurally similar to carbamazepine but has fewer drug interactions and fewer side effects, with the exception of an increased frequency of clinically significant hyponatraemia. Because it is chosen as the gentler option, the sodium is what nobody checks. Learn the direction, not a percentage: no source here quantifies incidence for either drug, so a confident figure is an invented one.

20.5 The rest of the causal field, and the sodium that is not really low

Psychotropics rarely act alone. The Maudsley notes that desmopressin, used for clozapine-induced enuresis, can itself cause hyponatraemia, and that antidepressants, anticonvulsants, valbenazine and many drugs for physical health conditions, including diuretics, angiotensin-converting enzyme inhibitors, angiotensin II receptor blockers and proton pump inhibitors, have been implicated, with risk probably additive across concomitant prescriptions, adding thiazides, opiates, NSAIDs, tramadol, omeprazole and trimethoprim. Reconcile the whole list, not only your part of it.

Non-drug causes come from the neurology side. Brazis records that although the syndrome is more often extrahypothalamic, partial damage to the supraoptic and paraventricular nuclei or neighbouring areas can produce it, through trauma, subarachnoid haemorrhage, hydrocephalus, tumours, meningitis, encephalitis or autoimmune disease such as anti-LGI1 antibodies, and that ectopic ADH production by a tumour, classically a carcinoma of the lung, does the same. His drug list independently names vincristine, chlorpropamide, cyclophosphamide, carbamazepine and chlorpromazine, corroborating the psychotropic causes from a neurological source.

Confirm the sodium is real before acting on it. **Pseudohyponatraemia** arises because severe hyperlipidaemia or hyperglycaemia produces a secondary increase in plasma volume, and both are commoner in people treated with antipsychotic drugs than in the general population, so both should be excluded as causes. The same antipsychotic that produces a genuine SIADH can manufacture a spurious low sodium. The metabolic criteria sit where this chapter deals with metabolic syndrome, obesity and mortality in serious mental illness.

20.6 Who is at risk, what to warn about, and how to monitor

The Maudsley's **high-risk factors for drug-induced hyponatraemia** are asked as straight recall.

- Older age and female sex, age perhaps the most important, so monitoring is essential in older people and especially older women.
- Major surgery.
- Previous hyponatraemia, or a low baseline sodium.
- Co-therapy with diuretics, NSAIDs, antipsychotics, carbamazepine, chemotherapy, calcium antagonists, angiotensin-converting enzyme inhibitors or laxatives.

- Reduced renal function, taken as a glomerular filtration rate below 50 mL/minute.
- Medical comorbidity such as hypothyroidism, diabetes, COPD, hypertension, head injury, cardiac failure, stroke and cancers, plus low body weight.

MNEMONIC

Old, One sex, Operations, Other drugs, Organ trouble, Once before

The six headings of the high-risk list. Low body weight is the seventh and sits outside the pattern.

Warning precedes testing. Everyone starting an antidepressant should be told about and watched for dizziness, nausea, lethargy, confusion, cramps and seizures, the gradient running from confusion, nausea, headache and lethargy in mild to moderate hyponatraemia to increasingly severe symptoms with seizures and coma as plasma sodium falls. Only Brazis attaches a number, at the severe end: anorexia, nausea, vomiting and irritability may progress to paranoid delusions and generalised seizures once serum sodium falls below **110 mEq/L**. The delusions matter: a patient with schizophrenia growing more paranoid on treatment reads as relapse, and increasing the antipsychotic makes everything worse.

For high-risk patients on antidepressants the schedule is serum sodium at baseline, at 2 and 4 weeks, then 3-monthly, with a caveat that should never be stripped: this frequency may not identify all cases, as serum sodium can drop precipitously. For antipsychotics no schedule is given, only that plasma sodium monitoring is desirable for everyone receiving them, particularly where several risk factors are present. The wider metabolic cadence belongs where this chapter deals with metabolic monitoring protocols and parameters for patients on antipsychotics.

20.7 Lithium: two sodium relationships running in opposite directions

Lithium is absent from the SIADH lists, and knowing why is worth a mark. Its characteristic lesion is the opposite one: lithium decreases the sensitivity of renal tubules to vasopressin and is a common cause of acquired **nephrogenic diabetes insipidus**, so fluid intake and urine output need close monitoring and patients risk hypernatraemia and severe dehydration where an impaired ability to drink to thirst meets an impaired ability to concentrate urine. SIADH retains water and drives sodium down; lithium loses water and drives sodium up. Lithium as a drug belongs to the Mood disorders: pharmacotherapy notes.

The second relationship runs the other way and reaches the ward. Any state inducing sodium depletion can cause lithium toxicity, because the relative paucity of sodium leads lithium to be preferentially reabsorbed by NHE3 in the proximal nephron. The scenario is specific: season, ambient temperature, sweating and strenuous exercise do not by themselves significantly alter lithium levels, and the common path is a patient depleted by gastrointestinal losses or heavy sweating who replaces them with free water rather than a balanced electrolyte solution. Here low sodium is cause and high lithium effect, the reverse of the arrow above; separating the two is what is tested.

IN INDIA

The sodium-depletion route to lithium toxicity is seasonal here in a way the source literature does not convey: outdoor work through the hot months, gastroenteritis through the wet ones, and the instinct to drink plain water afterwards assemble the described scenario exactly. A patient on lithium losing fluid to vomiting, diarrhoea or heavy sweating should rehydrate with oral rehydration solution, available in any chemist, rather than water alone. Monitoring diverges downward: the sodium schedule assumes repeat electrolytes are affordable and the patient returns to the same clinic, and where neither holds, the warning given to patient and family does more work than the calendar.

20.8 Management: LOCUS, FOCUS, MODUS

No dedicated national guideline covers this here; the Maudsley supplies the thresholds and rechallenge advice, Goodman and Gilman the drug options. **LOCUS** is set by the number. Mild hyponatraemia above 130 mmol/L may be manageable with fluid restriction, some suggesting increased sodium intake though this is likely impractical, with the antidepressant discontinued if symptoms persist: outpatient work. Above 125 mmol/L, monitor sodium daily until normal and withdraw the offending antidepressant as soon as possible, watching for headache, nausea, vomiting, muscle cramps, restlessness, lethargy, confusion and disorientation. Below **125 mmol/L**, refer urgently to specialist medical care and discontinue immediately, given the increased risk of life-threatening seizures, coma and respiratory arrest: emergency transfer, not a review appointment. The same passage warns that over-rapid correction may itself be harmful; the safe rate of rise and the demyelinating syndrome that follows breaching it are taught where this chapter deals with psychogenic polydipsia, the sodium-correction risk and mineral disturbance.

FOCUS shifts by phase: acutely the falling sodium and the offending drug together; short term, the daily sodium until back in range; mid term, whether antidepressant treatment resumes, since untreated depression is not a safe alternative; long term, the modifiable risk, most of which is co-prescription.

MODUS begins with subtraction: withdraw the offending drug and any additive co-prescription, then restrict fluid. Goodman and Gilman record that demeclocycline has been the preferred drug but that the **vasopressin V2 receptor antagonists** are now available: tolvaptan, a selective oral V2 antagonist FDA approved for clinically significant hypervolaemic and euvolaemic hyponatraemia, and conivaptan, a non-selective V1a and V2 antagonist FDA approved for hospitalised patients with euvolaemic and hypervolaemic hyponatraemia and available only for intravenous infusion. The Maudsley calls the vaptans aquaretics because they induce a highly hypotonic diuresis, with promise in hyponatraemia of various aetiologies including drug-related SIADH. No dose or duration for the three appears in this material, and no dosing regimen should come from memory. The Maudsley also lists lithium, cautioning that hyponatraemia predisposes to lithium toxicity, a poor choice in exactly the patient who would need it.

Restarting is the part candidates skip. **Rechallenge** carries real recurrence risk: hyponatraemia may recur with the same or a different SSRI but may be less likely with an antidepressant from another class. The strategy is to withdraw other drugs associated with hyponatraemia, since risk increases exponentially when antidepressants are combined with diuretics; prescribe from a different class, considering nortriptyline, lofepramine, mirtazapine or moclobemide, starting low and monitoring closely; consider water restriction and careful use of demeclocycline where it recurs and continued treatment is essential; and consider ketamine or esketamine, or ECT, where no standard antidepressant can be given. For antipsychotics the advice is deliberately weaker: switching is suggested, but data are insufficient to guide the choice and cross-sensitivity may occur, the individual perhaps being predisposed overall and the drug chosen unimportant. In a worked teaching case of depression in a medically ill older adult at high SIADH risk, where SSRIs and tricyclics are avoided, mirtazapine 15 mg at bedtime is preferred, to improve sleep and appetite and reduce nausea; read it as one case-illustrative starting dose, not a recommended regimen. The same logic applies where lung cancer with central nervous system metastasis already places SIADH risk unusually high.

PITFALL

One widely used text points the wrong way, and copying it would be dangerous. Kaplan and Sadock's Comprehensive Textbook of Psychiatry, in its eating-disorders section, states that whatever the aetiology fluid restriction is the treatment of choice, and that intranasal arginine vasopressin should be considered where fluid restriction does not sufficiently improve sodium. Every other source here runs the other way: the Maudsley lists desmopressin, given for clozapine-induced enuresis, among the causes, and the drugs named as treatment are vasopressin antagonists, the mechanistic opposite. A specialist nephrological manoeuvre does exist, in which desmopressin prevents over-rapid autocorrection while sodium is raised controllably, but that is not what the sentence conveys and is not a psychiatric procedure. Do not give desmopressin for hyponatraemia; refer for specialist medical care.

Serum osmolality is low in both mechanisms and only the urine discriminates; risk peaks in the first weeks and is gone by three to six months; oxcarbazepine is worse than carbamazepine; and below 125 mmol/L the drug stops today and the patient goes to a medical unit.

21 · Psychogenic polydipsia, the sodium-correction risk and mineral disturbance

Three questions sit inside this topic and examiners like all of them. One is behavioural, the drinking itself; one is iatrogenic and is the one that maims, a pons lost to the correction; one is the wider mineral field of calcium, magnesium, potassium and phosphate. One discipline ties them: establish the mechanism before acting on the number. A low sodium from drinking, one from inappropriate antidiuresis and one the body is defending on purpose behave differently, and two of the three call for doing less than the number invites. Inappropriate antidiuretic hormone secretion and the drug causes of hyponatraemia are treated separately in this chapter.

At least 20%	6 to 17%	Over 12 mEq/L/day	Under 10%
POLYDIPSIA AMONG CHRONIC PSYCHIATRIC INPATIENTS	WATER INTOXICATION, CHRONICALLY ILL HOSPITALISED PATIENTS	RISE AT WHICH PONTINE MYELIN BREAKS DOWN	PSYCHIATRIC SYMPTOMS IN HYPERPARATHYROIDISM

21.1 What primary polydipsia is, and what it is not

Primary polydipsia is increased water intake occurring in the absence of any impairment in water excretion. The kidney is working; the patient is drinking faster than a normal kidney can clear. That separates it from **secondary polydipsia**, where water is being lost obligatorily, the classic psychiatric example being the lithium-treated patient with nephrogenic diabetes insipidus replacing what the kidney throws away.

Both groups may show a low urine osmolality, and that shared finding is where candidates come unstuck. The discriminator is the serum: in lithium-driven secondary polydipsia sodium and osmolality stay normal because intake matches loss; in primary polydipsia both fall during binges, sometimes severely. Lithium as a drug and its non-thyroid organ-system adverse effects belong to the Mood disorders: pharmacotherapy notes.

TRAP **Assuming the lithium patient who drinks a great deal is therefore hyponatraemic.** Heavy drinking, dilute urine and lithium invite you to call it water intoxication. It is not: the drinking is compensatory, serum sodium and osmolality are normal, and fluid restriction dehydrates the patient and pushes the lithium level up. Read the serum, not the story.

The named variants appear in single-best-answer stems. Patients may drink from conditioned behaviour, which supplies the two eponymous presentations, **beer drinker's hyponatraemia** and **tea party epilepsy**, or from other psychogenic factors. A second group has hyperangiotensinaemia, described in renal patients who stay thirsty despite an electrolyte balance kept normal by haemodialysis. A third, rare group has hypothalamic disease, where the drinking compensates for a mild diabetes insipidus, the mirror image of the trap above.

21.2 How common, and in which population

Three different numbers get merged into one wrong number. They do not share a denominator.

CONSTRUCT	FIGURE	POPULATION IT IS BOUND TO
Excessive water intake, historical observation	approximately 25% of patients, noted by 1936	The severely mentally ill of the era, in whom it was called the commonest metabolic abnormality and was already linked to life-threatening hyponatraemia.
Polydipsia, the drinking behaviour, modern prevalence	at least 20%	Chronic psychiatric inpatients, from data gathered over 5 years in a state hospital. Note the direction: at least, not up to.
Water intoxication, the biochemical consequence	6 to 17%	Cross-sectional studies of chronically ill, hospitalised psychiatric patients.
Episodic hyponatraemia secondary to fluid overload	10%	A longitudinal study of severely ill patients with a diagnosis of schizophrenia.

One caution about independence: the modern figure in the clozapine literature and those in the prescribing literature rest substantially on the same body of work, so quoting both is not two confirmations. The ordering is the teaching point: a common behaviour, a smaller fraction of drinkers becoming biochemically water-intoxicated, a smaller group again cycling through it. That argues for screening, not for waiting.

21.3 Why they drink: two accounts that do not agree

The aetiology is poorly understood and the standard sources offer different accounts. The clozapine literature is unambiguous that the drinking is not due to excessive thirst but is motivated by delusions or by a psychic discomfort that the water binge relieves, making the behaviour closer to a delusionally motivated ritual than to a homeostatic drive.

The prescribing literature offers two drug-mediated hypotheses instead. The first is an extreme compensatory response to the anticholinergic adverse effects of some antipsychotic drugs, a dry mouth taken to its dangerous conclusion. The second is a receptor argument: post-synaptic dopamine receptor antagonism produces receptor supersensitivity, increased presynaptic dopamine release and elevated dopamine in the hypothalamus, driving thirst and polydipsia. Three observations support it: many reported cases occur in patients with long illness histories, many occur on antipsychotics with high D2 receptor affinity, and clozapine can improve polydipsia independently of improvement in psychosis.

That last observation cuts against the purely psychopathological account: if water intake falls while the delusions do not, the drinking was not simply a delusional act. Antipsychotic receptor pharmacology and drug selection belong to the Schizophrenia: management, antipsychotics and adverse effects notes. The correct answer names both accounts and settles neither.

21.4 Telling it apart at the bedside

Polydipsia with polyuria has a short differential a psychiatrist must run. It may reflect impaired renal concentrating capacity from deficiency of antidiuretic hormone, which is cranial diabetes

insipidus, or from failure of its action, which is nephrogenic diabetes insipidus. The nephrogenic form may be inherited or acquired, hypercalcaemia and hypokalaemia both impairing that action: a calcium or potassium disturbance can itself manufacture a polyuria that looks like drinking. Or it may simply follow vastly excessive intake.

The bedside sign is **nocturnal polyuria**, generally not a feature of psychogenic polydipsia because the drinking stops when the patient sleeps, whereas a kidney that cannot concentrate keeps producing urine through the night. Hutchison's own hedge must travel with the sign: this is not an absolute distinction, and further investigation of urine concentrating capacity is usually required.

PITFALL

Sodium in water intoxication shows diurnal variation, dropping as the day progresses because the drinking accumulates across waking hours. A morning-round sample is drawn at the patient's best moment, and a ward that samples only then will keep recording acceptable sodium in a patient who is seizing by evening. Where the concern is water intoxication rather than a drug effect, time the sample to the risk.

Eating disorders are a second, quite separate population. Hyponatraemia occurs commonly in bulimia nervosa and often in anorexia nervosa, from increased water intake diluting the serum sodium or from a syndrome of inappropriate antidiuretic hormone secretion. The intake is deliberate: patients may drink to suppress appetite and support restriction, or, where treatment involves weigh-ins, practise **water loading** to inflate the recorded weight. Because hyponatraemia increases seizure risk it must be addressed quickly, and whatever its aetiology here, fluid restriction is the treatment of choice.

21.5 Management of polydipsia and water intoxication

LOCUS. Managed where the patient already is, on the ward or in the outpatient clinic, because the behaviour is chronic and the monitoring longitudinal. The escalation point is explicit: refer urgently to specialist medical care if the serum sodium falls below **125 mmol/L**. Below that line the psychiatrist's job is handing over an accurate account of the mechanism rather than managing the electrolyte.

FOCUS. Acutely, the water intoxication and the seizure risk that rides with it, plus anything exogenous making the sodium worse. Longer term, the drinking behaviour, the only durable fix.

MODUS. Three evidence-based options are named for preventing water intoxication: **targeted fluid restriction**, clozapine therapy, and removal of agents that may themselves be causing hyponatraemia, the examples given being carbamazepine, valproate and sodium-wasting diuretics. Fluid restriction is coupled to careful sodium monitoring, which must take the diurnal variation above into account. Note what the sources do not supply: none states a volume, so no litre-per-day target belongs in the plan.

Clozapine has demonstrated efficacy in schizophrenia with psychogenic polydipsia at doses as low as **300 mg per day**. The evidence base is thin and should be quoted as thin: a 24-week open-

label study of eight male patients with schizophrenia and polydipsia who had already met treatment-resistance criteria, run as sequential 6-week blocks of clozapine at 300, then 600, then 900 mg per day where tolerated. On typical antipsychotics serum and urine osmolality stayed grossly abnormal; on clozapine mean plasma osmolality normalised, rising by 15.2 mosm/kg (95% CI 5.5 to 25.0), the effect evident at the lowest dose step. There is no polydipsia-specific titration to learn: it is that of treatment-resistant schizophrenia. The prescribing literature agrees clozapine raises plasma and urine osmolality, while cautioning that this is consistent with reduced fluid intake and not clearly related to mental state.

Two negatives are as examinable as the positives. There is no evidence that reducing or increasing the antipsychotic dose improves serum sodium in water-intoxicated patients, so dose manipulation is not a treatment. And a long list of other agents has been tried with limited data behind any, which is reason to distrust a confidently named second-line drug.

TRAP **Reaching for demeclocycline because a guideline mentions it.** It is included in some practice guidelines for psychogenic polydipsia. Against that, its mechanism is to interfere with antidiuretic hormone and increase water excretion, already running at capacity in a pure water-intoxication patient, so the rationale is debatable and some published cases may have been complicated by undiagnosed inappropriate antidiuresis. A single small randomised trial showed no benefit. The same source nonetheless keeps it for inappropriate antidiuresis, recording the conflict without resolving it. The question is whether you established the mechanism first, and urine osmolality is what establishes it.

21.6 The correction, not the illness, is what maims

Overly rapid correction of hyponatraemia can produce central pontine myelinolysis, the prime example of the **osmotic demyelination syndrome**. Vulnerable patients lose myelin in the pons and are often left with permanent nystagmus, quadriparesis and ataxia. Three groups are named repeatedly: malnourished alcoholics, liver transplant recipients, and pregnant women with hyperemesis gravidarum, alongside others prone to hyponatraemia. The prescribing literature calls it irreversible.

> The pons is destroyed by the correction and not by the hyponatraemia, and it is chronic hyponatraemia whose correction destroys it.

The timing is what makes central pontine myelinolysis a psychiatric problem. It occurs days after the overly rapid correction of chronic hyponatraemia, so the patient improves first and deteriorates afterwards, by which point the causal link is easy to miss. It typically presents with delirium and a flaccid quadriparesis, and uncommonly with delirium alone, and that second version is the one a psychiatrist is asked to see. In alcohol-dependent patients the picture is quadriplegia with pseudobulbar palsy. Delirium as a syndrome belongs to the Stroke, TBI, delirium and focal brain syndromes notes; what belongs here is that a new delirium in a patient whose sodium was corrected days earlier is a structural event until proven otherwise.

RULE **In hyponatraemia, every decision is a decision about rate.** The single number the corpus supplies for the rate of sodium correction is a threshold for harm: correction faster than **12 mEq/L per day** is what leads to myelin breakdown in the pons. Read the direction carefully: that figure says where damage occurs, not where treatment should aim, and must not be converted into a target. No recommended maximum safe rate is stated anywhere in the standard psychiatric references, and a rate quoted from memory is a rate you cannot attribute.

21.7 The other end of the range, and a sodium that should be left alone

Essential hypernatraemia is a rare syndrome of decreased antidiuretic hormone release with absence of thirst, diagnosed on four criteria: hypernatraemia without a corresponding fluid deficiency, preserved renal responsiveness to the hormone, impaired secretion in the face of hypernatraemia, and absence of thirst despite preserved conscious behaviour. Lesions have affected the tuberal region or the whole hypothalamus. Some patients tolerate a high sodium well enough to develop water intoxication when treated, an iatrogenic hazard that runs opposite to central pontine myelinolysis. Sodium reaching 170 mEq/L here is attended by muscle cramping, tenderness and weakness, fever, anorexia, paranoia and lethargy. That figure is bound to this syndrome and its unusual tolerance; generalising it into a symptom threshold for hypernatraemia at large misuses it. No source in the standard references defines hypernatraemia by a serum sodium value.

Reset osmostat hyponatraemia is the third mechanism and the one most likely to be mismanaged. Its criteria are normovolaemic hypotonic hyponatraemia; normal renal, adrenal and thyroid function; the ability to concentrate urine when serum tonicity is raised above the reset level; the ability to excrete a water load; and maintenance of normal sodium balance without correction of the hyponatraemia during salt loading. The last pair does the work: preserved excretion of a water load separates it from inappropriate antidiuresis, and the failure of salt loading says the body is defending a lower set point rather than failing to reach a normal one. This is the stable, asymptomatic, treatment-resistant low sodium that keeps prompting intervention, and precisely the patient in whom chasing the number carries the demyelination risk for no gain.

21.8 Calcium: the mineral with a psychiatric reputation it does not quite deserve

Hypercalcaemia is the classic examinable mineral cause of psychiatric presentation, and it is not a common one. Hyperparathyroidism is usually diagnosed after an incidental hypercalcaemia on routine bloods, and 50% of patients are asymptomatic. Where symptoms occur they are principally those of the hypercalcaemia itself: fatigue, malaise, proximal muscle weakness, renal colic, abdominal pain and cognitive decline. Psychiatric symptoms belong to the classic tetrad, but their prevalence is likely to be **less than 10%**, and they correlate with both the duration and the severity of the hypercalcaemia.

MNEMONIC

Bones, stones, moans and psychic groans

The symptomatic tetrad of hyperparathyroidism. Worth holding for the viva with the qualifier attached: the psychic groans are the least frequent of the four, not the headline.

The psychiatric profile is reasonably specific. Affective disorders, predominantly depressive, are the most frequently recorded, with anxiety often comorbid; psychotic disorders are reported but infrequent, with persecutory and paranoid delusions predominating; cognitive changes usually take the form of short-term memory loss, disorders of attention and acute confusional states; and severe derangement of calcium metabolism may present with somnolence and coma. Neuropsychiatric symptoms are often the initial presentation in older adults or those with limited cognitive reserve, the group most likely to be read as dementia. Symptoms generally respond to treatment such as parathyroidectomy. The severity correlation is stated repeatedly, but no calcium value is ever attached to it.

IN INDIA

Two practical points bite here. Where parathyroid hormone is mildly raised against a normal calcium, most commonly from vitamin D insufficiency, psychiatric disturbance is uncommon. Given how often that insufficiency is met in practice here, the combination is a finding to note, not an explanation for the presentation, and treating it as the cause closes an inquiry that should stay open. Second, magnesium is not on the routine electrolyte panel in most laboratories and must be requested by name at additional cost. Given how much of what follows turns on it, a psychiatrist who does not ask will not get it.

At the other end, hypocalcaemia announces itself through neuromuscular excitability: paraesthesiae, muscle cramps, carpopedal spasm and facial grimacing progressing to laryngeal spasm and convulsions. Examination may reveal tetany, reduced or absent deep tendon reflexes, papilloedema, and QT prolongation, which matters directly before any QT-lengthening psychotropic is added. The elicited sign is **Chvostek's sign**: gentle percussion over the proximal facial nerve as it exits the parotid gland is positive if it evokes involuntary facial twitching.

Delirium is now understood to be the commonest neuropsychiatric manifestation of hypocalcaemia. One large study reported cognitive impairment in 39%, non-specific affective disturbance in 21%, affective or neurotic symptoms in 12% and psychotic symptoms in 11%. Those categories are not mutually exclusive and belong to that study's population, so quote them as a profile rather than summing them. Severity again tracks the degree of hypocalcaemia and normalisation resolves the symptoms, with one exception.

21.9 Magnesium, potassium and the corrections that will not hold

RULE **Correct the magnesium, or the correction you are proud of will not hold.** This appears three times in three separate contexts across the standard references. Persistent psychosis has been recorded in hypocalcaemia where coexisting **hypomagnesaemia** was not addressed. Refractory cramping after apparently adequate correction of hypocalcaemia can be due to an associated hypomagnesaemia. And hypomagnesaemia, which may co-occur with hypokalaemia, must also be addressed because it inhibits prolonged normalisation of potassium. Whenever a corrected electrolyte drifts back down or a corrected patient fails to improve, magnesium is the first thing to check.

Hypokalaemia is the disturbance a psychiatrist is most likely to meet outside the ward, and it is dangerous rather than merely abnormal: it can cause cardiac arrhythmias and occurs commonly in individuals with recurrent vomiting. Potassium should be monitored regularly where purging is self-reported, with electrocardiograms for arrhythmias and QT and QTc prolongation. The caveat that is easy to lose: that prolongation may occur in eating disorders even without any electrolyte abnormality, so a normal potassium does not license skipping the tracing. In the opposite direction, hyperkalaemia tends to produce bradyarrhythmias and heart block rather than the lethal tachyarrhythmias of hypokalaemia.

The eating-disorder risk assessments give the only numeric mineral thresholds here. Among the blood-test markers of high physical risk are **hypokalaemia below 3.0 mmol/L** and hypophosphataemia below 0.5 mmol/L, alongside a very low body mass index, hypothermia, bradycardia, hypotension, postural hypotension, arrhythmia, hypoglycaemia and neutropenia. Both values are internally plausible, but the table they come from cannot be fully attributed, so treat them as indicative rather than quotable. Beyond them the standard references supply no reference range for serum calcium, ionised calcium, magnesium or potassium, which is why this material is examined as phenomenology and mechanism rather than numbers.

Refeeding is where all of this becomes predictable rather than incidental. Electrolytes should be monitored during refeeding even when normal at initial presentation, because hypophosphataemia and hypomagnesaemia may develop during the early phases; liver enzymes are usually normal but may increase transiently. No numeric threshold or calorie-escalation schedule for refeeding is given in these sources, and neither should be improvised.

GOOD TO REMEMBER

Intermittent cramps and bilateral pins and needles can be due to a decreased level of circulating ionised calcium, and that fall does not require hypoparathyroidism. Any rise in extracellular pH lowers the ionised fraction, and the best recognised setting is the respiratory alkalosis of a hyperventilatory state, with hypokalaemia producing the same effect through a metabolic alkalosis. That is the mechanism behind the perioral and peripheral paraesthesiae and the carpopedal spasm of a panic attack, worth explaining to the patient because the phenomenon is frightening and the explanation itself therapeutic. One rider: paraesthesiae confined to the hand may be median nerve compression at the wrist.

Hold the whole mineral set together as one line on the delirium screen. The metabolic causes of delirium listed in the standard texts include hyponatraemia and hypernatraemia, hypoglycaemia and hyperglycaemia, hypocalcaemia and hypercalcaemia, and hypomagnesaemia and hypermagnesaemia, with the general medical lists adding liver dysfunction, uraemia, thyroid dysfunction and nutritional deficiency, particularly of vitamin B12, thiamine and niacin.

- Sodium: low from drinking, from inappropriate antidiuresis or from a reset set point; the mechanism decides the response.
- Calcium: high produces depression and cognitive change; low produces tetany, seizures and delirium.
- Magnesium: rarely the presenting abnormality, frequently the reason another correction fails.
- Potassium: low in the purging patient, arrhythmogenic, not normalisable while magnesium is low.
- Phosphate: the refeeding mineral, falling early even when admission bloods were clean.

Encephalopathies, reproductive endocrinology and toxins

22 · Hepatic encephalopathy: grading, pathophysiology and psychiatric management implications

The psychiatrist usually meets this condition before the physician names it. The referral arrives as a behavioural problem: a man with cirrhosis irritable and muddled overnight, or a woman on a de-addiction ward whose attention has narrowed over a week. Kaufman's Clinical Neurology for Psychiatrists sets out the sequence. With hepatic insufficiency, mental function and consciousness steadily decline, and mild confusion with either lethargy or, less frequently, agitation may precede coma and overtly abnormal liver function tests. That last clause inverts the order most trainees expect.

Marks concentrate here for a reason only partly about the disease. Almost every specific a candidate confidently recalls, the numbered grades, the ammonia cut-off defining each one, the millilitres of lactulose, is either unsourced or contradicted by the standard references. Delirium as a syndrome belongs to the Stroke, TBI, delirium and focal brain syndromes notes; choosing and dosing a psychotropic in liver failure belongs where this chapter deals with prescribing in hepatic and renal impairment.

22.1 The disease, its nosological place and the full constellation

Kaplan and Sadock's Comprehensive Textbook of Psychiatry places hepatic encephalopathy as a **delirium due to another medical condition**. That places the disease; it does not license re-deriving the delirium construct. The corollary matters: fluctuation, inattention and altered arousal are expected, and a stable, well-oriented presentation should prompt another explanation.

The picture is far wider than confusion, which is why it lands on a psychiatric desk. The same source describes a condition commonly encountered in significant liver disease, presenting as a constellation: altered consciousness up to stupor or coma, cognitive impairment, confusion or disorientation, affective and emotional dysregulation, psychosis, behavioural and bioregulatory disturbance, and motor signs including asterixis, tremor, brisk reflexes, increased tone, ataxic gait, bradykinesia, slurred speech and incoordination. An apparent first psychotic episode, a manic disinhibition and a personality change in liver disease all sit on that list, and none defers the metabolic assessment.

RULE **Normal liver function tests do not exclude hepatic encephalopathy.** The mental state declines before the biochemistry does, so the tempting sequence, abnormal tests first and encephalopathy second, runs backwards. This is a clinical diagnosis supported by the laboratory, not gated by it.

22.2 The named forms, and the grading scheme these references do not carry

Honesty about what these references contain is what this section offers. They name forms consistently, minimal, subclinical, overt, fulminant and persistent, and never attach a numbered grade to any of those words. The numbered severity ladder in wide hepatological use is not reproduced here, because a grading scheme recalled from memory is exactly the specific that reaches a resident wrong. Take it from a hepatology source that prints it.

What they carry instead is prognostic direction. Kaplan and Sadock state that **minimal hepatic encephalopathy** predicts overt encephalopathy, and that the fulminant form is associated with intracranial hypertension, cerebral oedema, seizures and death before transplantation. Those statements bracket the condition: an abnormality detectable only on testing is a forecast, and at the other end the syndrome stops being psychiatric at all.

Reversibility ties the forms together and is why this is not a dementia. **Persistent hepatic encephalopathy** is rare, seen predominantly with extensive portacaval collateral circulation or after surgical or transjugular portosystemic stent shunting. Persistence is structural: where portal blood has a permanent route around the liver, the insult does not stop. The dementia syndrome and its workup belong to the Dementias and neurocognitive disorders notes; what belongs here is that cognitive impairment in liver disease carries a presumption of reversibility.

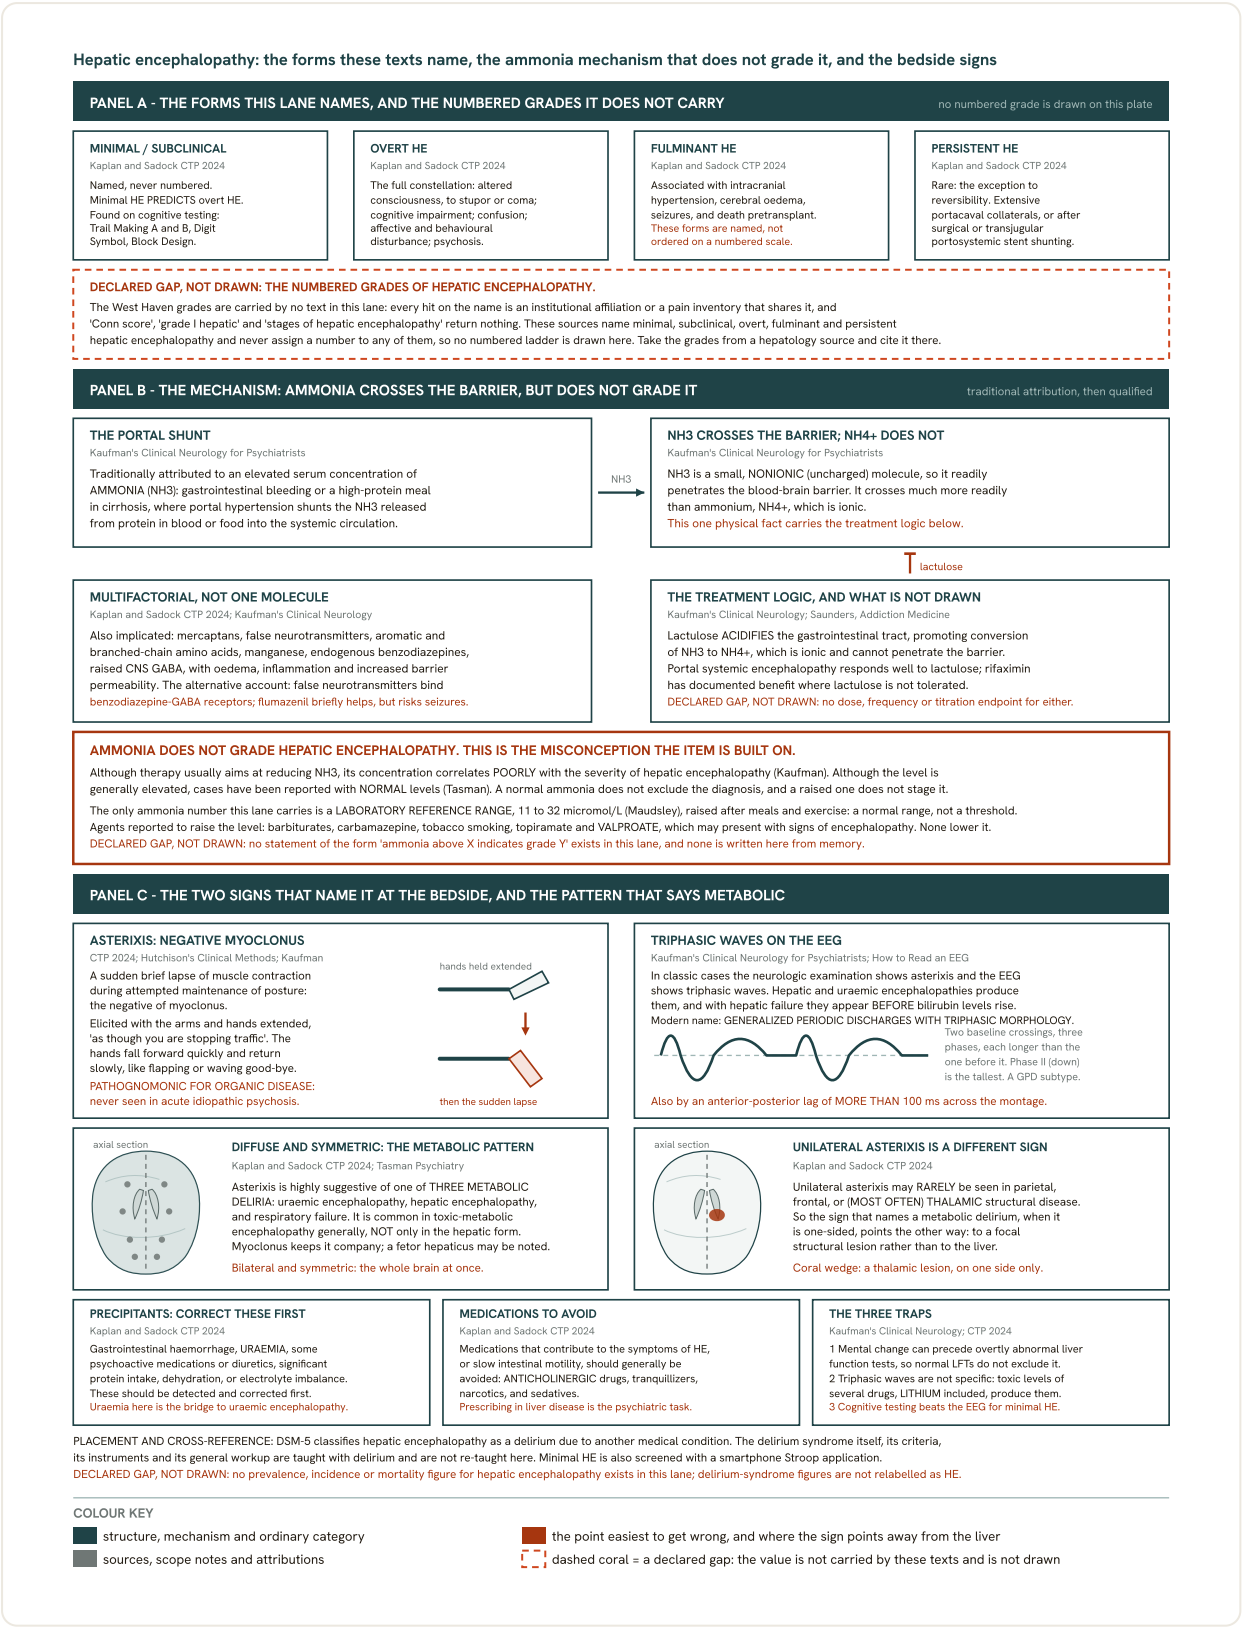


Fig 33.9 Hepatic encephalopathy: the forms these texts name, the ammonia mechanism, and the two bedside signs. No text in this lane numbers the severity: minimal, subclinical, overt, fulminant and persistent are named and none is graded, so the plate draws the vocabulary and marks the numbered ladder as a declared gap. The mechanism survives with its limit: ammonia crosses the barrier because it is small and uncharged, and lactulose converts it to ammonium, yet its concentration correlates poorly with severity and may be normal. Asterixis is pathognomonic for organic disease but not specific to the liver, and when unilateral points at structural disease instead. (Kaufman's Clinical Neurology for Psychiatrists; Kaplan and Sadock 2024; Tasman; Hutchison's Clinical Methods; How to Read an EEG; Maudsley Prescribing Guidelines; Saunders, Addiction Medicine.)

The accompanying figure sets the severity ladder alongside the mechanism and the two anchoring signs.

FORM NAMED IN THE SOURCES	WHAT THE SOURCES STATE	WHY IT MATTERS
Minimal	Predicts overt encephalopathy	Detectable only on testing; not benign
Subclinical and overt	Named; given no numbered grade in these texts	See a hepatology source
Fulminant	Intracranial hypertension, cerebral oedema, seizures, death before transplantation	A transplant and intensive-care problem
Persistent	Rare; portacaval collaterals or post-shunting	The structural exception to reversibility

22.3 The ammonia route, and why one molecule's charge decides everything

The traditional account drives the treatment. Kaufman records the classical attribution to a raised serum ammonia, operating in the common scenario of gastrointestinal bleeding or a high-protein meal in cirrhosis, where cirrhosis-induced portal hypertension shunts ammonia released from protein in blood or food straight into the systemic circulation. The mechanism is as much anatomical as biochemical: the nitrogen load is ordinary, but it bypasses a liver that would have cleared it. A gastrointestinal bleed is a reliable precipitant because it is a high-protein meal in the gut lumen.

Access to the brain is decided by charge alone. Ammonia is small, non-ionic and uncharged, so it penetrates the blood-brain barrier readily and far more readily than the ionic form, **ammonium**. That is the hinge of the topic: shift the equilibrium between the two species and you change how much nitrogen reaches the brain.

This is what the first-line treatment does. **Lactulose** is given to acidify the gastrointestinal tract, promoting conversion of ammonia to ammonium, which is ionic and cannot penetrate the blood-brain barrier. It is not a systemic scavenger; it traps nitrogen in the gut lumen in a form that cannot follow the portal shunt. That guards against calling lactulose a purgative: catharsis contributes, but the stated mechanism is a pH-driven speciation shift.

22.4 Ammonia is a treatment target and not a severity marker

Here the standard references genuinely disagree, and the disagreement is worth teaching rather than flattening. Kaplan and Sadock write that although the pathogenesis is multifactorial, the predominant treatment strategy is to reduce ammonia production from intestinal protein and reduce its absorption. Read alone, that invites treating the ammonia level as the disease. Two sources block the inference. Kaufman states that although therapy aims at lowering ammonia, its concentration correlates poorly with the severity of hepatic encephalopathy, and prints the opposite proposition as the wrong answer to a review question. Tasman's Psychiatry closes the loop: although the level is generally elevated, cases have been reported with normal levels.

Ammonia is a valid therapeutic target and an invalid severity marker: lowering it helps, while measuring it tells you neither how ill the patient is nor, when normal, that the patient is well. No numeric threshold for diagnosing, staging or grading appears in these references. The one figure available is a laboratory reference range from The Maudsley Prescribing Guidelines, **11-32 micromol/L**, rising after meals and exercise: a normal range for the assay, not a threshold for the disease.

One entry in that same table is squarely a psychiatric problem. The Maudsley lists the agents reported to raise ammonia, barbiturates, carbamazepine, tobacco smoking, topiramate and valproate, noting that valproate may present with signs of encephalopathy, and records none known to lower it. Valproate-induced hyperammonaemic encephalopathy is the one situation in which a psychiatrist sends an ammonia level on a structurally normal liver; the prescribing decisions belong to the Mood disorders: pharmacotherapy notes.

TRAP **Writing that ammonia levels are used to grade hepatic encephalopathy.** No source supports it and two contradict it. The trap is built on a true premise, that ammonia is the main treatment target. Learn the two negatives instead: the concentration correlates poorly with severity, and cases with normal levels have been reported.

Ammonia is a treatment target, never a severity marker: a normal level does not exclude hepatic encephalopathy, a raised level does not grade it, and the only ammonia figure worth carrying is a laboratory reference range.

22.5 The multifactorial account and the GABA hypothesis

Ammonia tracks severity poorly because it is one contributor among many. Kaplan and Sadock implicate mercaptans, false neurotransmitters, the aromatic amino acids phenylalanine, tyrosine and tryptophan, the branched-chain amino acids, manganese, endogenous benzodiazepines and raised central gamma-aminobutyric acid, alongside oedema, inflammation and increased blood-brain barrier permeability. The list is dominated by neurotransmitter precursors and neuromodulators rather than toxins, which is why the picture is behavioural before it is neurological, and its last three items are processes rather than metabolites, which is why a patient deteriorates at an unchanged ammonia level when sepsis or dehydration supervenes.

The best-developed alternative is testable. Kaufman describes an account emphasising **false neurotransmitters** that bind benzodiazepine-gamma-aminobutyric acid receptors and increase GABA activity, and notes that **flumazenil**, a benzodiazepine antagonist acting at those receptors, temporarily reduces symptoms though it carries some risk of precipitating seizures. Hold on to the word temporarily: this is a mechanistic demonstration, not a strategy, and the seizure risk is stated in the same breath. No dose appears in these references and none is offered here.

GOOD TO REMEMBER

The endogenous benzodiazepine account has a prescribing consequence easily lost on a busy ward. If part of the encephalopathy is driven by activity at the benzodiazepine-GABA receptor complex, a benzodiazepine given for the agitation adds to the mechanism producing it. That is the reason behind the blunt instruction to avoid sedatives.

22.6 Asterixis: what it is, how to elicit it, and what it does not mean

The first of the two bedside signs is a movement disorder defined by absence. Kaplan and Sadock build it from its positive counterpart: myoclonus is a sudden, jerky, shock-like movement, more discontinuous than chorea or tremor, and its negative is **asterixis**, a sudden lapse of muscle contraction during attempted maintenance of posture. Hutchison's Clinical Methods records a negative myoclonus, a sudden brief loss of tone, particularly associated with hepatic encephalopathy. The term is not pedantry: the sign exists only while a posture is actively held, which dictates how it is sought.

Kaufman describes the manoeuvre. Asterixis is elicited by having the patient extend the arms and hands as though stopping traffic; the hands then intermittently and quickly fall forward and slowly return, as though flapping or waving goodbye. The asymmetry is the diagnostic detail: a fast drop with slow recovery is a lapse of tone, a symmetrical oscillation a tremor. Kaplan and Sadock add that it should be carefully sought, because it is pathognomonic for organic disease and is never seen in acute idiopathic psychosis or other non-organic disorders, and Addiction medicine, listing the findings in alcohol-related liver disease, says to check for asterixis if hepatic encephalopathy is suspected. Tasman adds the accompanying olfactory sign, a characteristic bad breath or **fetor hepaticus**.

The trap runs the other way. Kaplan and Sadock note that both phenomena, but more sensitively asterixis, are common in toxic-metabolic encephalopathy and not just in hepatic encephalopathy. Tasman narrows it: asterixis is highly suggestive of one of three metabolic deliria, namely uraemic encephalopathy, hepatic encephalopathy and respiratory failure. One refinement is counter-intuitive. Unilateral asterixis may rarely be seen in parietal, frontal or, most often, thalamic structural disease. A flap on one side only is a localising sign, not a metabolic one.

MNEMONIC

Liver, kidney, lung: the three organs that flap

Asterixis in a delirious patient points at **liver**, **kidney** or **lung**. Then ask the question that is not an organ: is the flap bilateral? If one-sided, think structurally, thalamus first.

TRAP **Reading asterixis as specific to liver failure.** Candidates learn the sign attached to hepatic encephalopathy and use it as though it named the organ. The safe formulation has two halves: it is pathognomonic for organic disease, but within organic disease it is a metabolic pointer with more than one destination.

22.7 The electroencephalogram, the triphasic wave and the cognitive tests

The second sign is electrographic, and Kaufman's classic case pairs the two: the examination demonstrates asterixis and the electroencephalogram shows triphasic waves. Most of what a metabolic encephalopathy does to an EEG is non-specific. Background slowing and disorganisation are typical but shared by numerous conditions; a few findings do suggest an aetiology, and **triphasic waves** often indicate hepatic or uraemic encephalopathy. In hepatic failure, triphasic waves usually appear before bilirubin levels rise, the electrographic version of the rule that mental change precedes the laboratory. And the finding may also be seen with toxic levels of several medications, including lithium: triphasic waves in a lithium-treated patient are not automatically hepatic, and that side of the question belongs to the Mood disorders: pharmacotherapy notes.

The modern term is worth a mark. How to Read an EEG classifies it as **generalized periodic discharges with triphasic morphology**, a common subtype seen in toxic and metabolic encephalopathies such as the hepatic and uraemic ones, but also with hypoxic and epileptic encephalopathies, and less commonly as the electrographic manifestation of non-convulsive status epilepticus. That last possibility is why morphology is read rather than pattern-matched. The same source cautions that no feature reliably separates ictal from non-ictal triphasic waves, and that on one retrospective study almost one-third of metabolic triphasic waves may show a definitive response to medication.

ELECTROGRAPHIC FEATURE	WHAT IT CONTRIBUTES
Two baseline crossings, three lengthening phases, tallest middle positive phase	Establishes triphasic morphology
Anterior-posterior lag of more than 100 ms	Supports the triphasic identification
Overriding spikes or sharp waves	Suggests an ictal pattern
Frequencies faster than 2.5 Hz	Suggests an ictal pattern
No change with state change or stimulation	Suggests an ictal pattern

At the mild end, testing wins. Kaplan and Sadock hold that although the electroencephalogram may show generalised slowing of the dominant rhythm, cognitive testing is more efficient and perhaps more sensitive to minimal hepatic encephalopathy. They name the Trail Making Tests A and B, Digit Symbol and Block Design, useful for minimal disease but not commonly employed because of time and expertise limitations, and a smartphone application using the Stroop cognitive screening test, validated against a standard neurocognitive battery and possibly a more feasible way to detect and track it. The instruments and their norms belong to the Dementias and neurocognitive disorders notes.

11-32
MICROMOL/L, LAB RANGE, NOT A THRESHOLD

>100 ms
ANTERIOR-POSTERIOR LAG, TRIPHASIC WAVE

>2.5 Hz
FASTER RUNS SUGGEST ICTAL

ALMOST ONE-THIRD
METABOLIC TRIPHASICS MAY RESPOND TO MEDICATION, ONE STUDY

IN INDIA

Neither the formal neuropsychological battery nor the electroencephalogram is realistically available where most of this work happens. District psychiatry units and de-addiction services see decompensated alcohol-related liver disease constantly, without same-week access to an EEG technician or a neuropsychologist, and the paper tests above are already described as not commonly employed for want of time and expertise. So the bedside asterixis check carries disproportionate weight in every confused patient with liver disease: it is free and, unlike the EEG, available at midnight. A validated smartphone-based screen is the one feasible route to tracking the minimal form.

22.8 Management implications for the psychiatrist: locus, focus and modus

LOCUS. Care happens in three places, all of them the psychiatrist's. The emergency and acute medical locus owns the fulminant form. The general hospital ward and the consultation-liaison referral own the common presentation, where the psychiatrist is asked to manage behaviour and should instead sequence an aetiological search. The outpatient and de-addiction locus owns the minimal form and relapse prevention. The hepatology society guidance governing the medical arm is not carried by the psychiatric references drawn on here: the positions below come from Kaplan and Sadock's Comprehensive Textbook of Psychiatry, Kaufman's Clinical Neurology for Psychiatrists and the addiction medicine literature.

FOCUS. The first target is not the mental state but the precipitant, and the instruction is explicit that these problems should be detected and corrected first:

- Gastrointestinal haemorrhage.
- Uraemia.
- Use of some psychoactive medications or diuretics.
- Significant protein intake.
- Dehydration or electrolyte imbalance.

Two of those are iatrogenic, and uraemia is the bridge to the other systemic encephalopathies dealt with elsewhere in this chapter. The short-term target is the encephalopathy once the precipitant is addressed; the mid-term target is recurrence, meaning the drinking, the diuretics and the nutritional state; the long-term target is the transition to overt disease.

MODUS. The gut-directed drug modes are the mainstay. Addiction medicine records that portal systemic encephalopathy responds well to lactulose, with **rifaximin** now increasingly available and of documented benefit where lactulose is not tolerated. No dose, frequency or titration endpoint for either agent appears in any reference used here, and none is invented: know the mechanism and the order of preference, and take the regimen from a prescribing source. The psychiatric mode is largely subtractive. Medications that contribute to the symptoms of hepatic encephalopathy or slow intestinal motility, namely anticholinergic drugs, tranquillisers, narcotics and sedatives, should generally be avoided. That covers most of what a liaison psychiatrist is tempted to prescribe here; which psychotropic may then be used is dealt with where this chapter

covers prescribing in hepatic and renal impairment. The procedural mode matters for its shadow: shunting is where the rare persistent form appears, so an encephalopathy that does not clear deserves a question about shunt anatomy. The social mode is the most decisive, because abstinence, nutrition and adherence to a gut-directed regimen determine whether any of the above holds.

PITFALL

Sedating the agitated, disinhibited cirrhotic patient because the presentation looks behavioural. The reasoning is ordinary: the patient is disturbed, the referral asked for sedation, the drug settles people. It does three unhelpful things at once. It adds to the receptor activity implicated in the mechanism, it slows intestinal motility in a patient whose treatment depends on gut transit, and it obscures the conscious level that is the main index of improvement.

23 · Uraemic, hypoxic and other systemic encephalopathies

These are the encephalopathies in which the failing organ is never the brain, and the psychiatrist is called first. A dialysis patient stops sleeping, turns irritable, then frankly paranoid. A patient with advancing neuromuscular weakness becomes overwhelmingly anxious overnight. A young man resuscitated after a hanging attempt opens his eyes and says nothing that means anything. In each the reflex psychiatric response is the wrong one.

Hepatic encephalopathy and its grading belong where this chapter deals with hepatic encephalopathy; dose adjustment in renal failure belongs where this chapter deals with prescribing psychotropics in hepatic and renal impairment. Delirium as a syndrome belongs to the Stroke, TBI, delirium and focal brain syndromes notes; the dementia syndrome and its investigation belong to the Dementias and neurocognitive disorders notes.

23.1 Why uraemia poisons the brain, and why no single molecule is to blame

Kaplan and Sadock's Comprehensive Textbook of Psychiatry frames **uraemic encephalopathy** as the endpoint of broad systemic failure rather than of one toxin. Chronic renal failure frequently produces multiple pathophysiologic changes across catabolic, metabolic and endocrinologic processes, and the accumulation of organic substances acting as neurotoxins is believed to drive the encephalopathy. The substances named are urea, uric acid, guanidine compounds, hippuric acid, various amino acids and indoleacetate among others, but no single metabolite has been identified as the sole cause.

That last clause is the examination point: as with ammonia in hepatic encephalopathy, the marker everyone measures explains none of the mental state.

A second set of contributors is the correctable one: hormonal elevations, electrolyte imbalances including acidosis, hyponatraemia, hyperkalaemia, hypocalcaemia and hypermagnesaemia, anaemia, malnutrition, and central nervous system factors comprising increased calcium with decreased GABA and glycine activity. That is a management list disguised as a pathophysiology list: the neurotoxins leave only with dialysis or a new kidney, while the electrolytes, the anaemia and the nutritional state can be addressed tonight.

RULE **In a systemic encephalopathy the brain is the readout, not the lesion.** Imaging excludes a structural competitor rather than finding the cause; the electroencephalogram grades severity but does not name the organ; the treatment that works is directed at the systemic derangement. A psychiatrist who answers uraemic paranoia with an antipsychotic before asking when the patient last dialysed has answered the wrong question.

23.2 The clinical picture, from insomnia to coma

The fluctuation is what identifies the disease as metabolic. Symptoms typically fluctuate and can begin insidiously, with patients reporting mild irritability, insomnia, or cognitive, attentional and memory impairment. Patients may complain of slurred speech, muscle twitches or restless legs. Symptoms can then progress, slowly or rapidly, to confusion, lethargy, seizures, overt delirium, psychosis, catatonia, and stupor or coma.

Three features matter. Irritability, insomnia and inattention in chronic kidney disease get referred as depression or non-compliance; the discriminator is the attentional deficit, which does not belong to a mood disorder. Psychosis and catatonia are explicitly on the list, so a catatonic presentation in renal failure is not a coexisting psychiatric diagnosis. And restless legs are routinely misread as akathisia once the patient is on an antipsychotic.

Kaufman's Clinical Neurology for Psychiatrists records that renal insufficiency is a common cause of neuropathy and occurs almost universally in patients on maintenance haemodialysis, so stocking-glove sensory loss is expected. The negative that follows is more useful: cancer chemotherapy agents and antibiotics, including those used for tuberculosis and HIV disease, routinely cause neuropathy, whereas antipsychotics, antidepressants and antiepileptic drugs, with the exception of phenytoin, do not.

23.3 Asterixis and myoclonus: the two signs that localise to an organ

Two involuntary movements do most of the bedside work. Tasman's Psychiatry states that **asterixis** is highly suggestive of one of three metabolic deliria: uraemic encephalopathy, hepatic encephalopathy and respiratory failure.

MNEMONIC

Liver, Kidney, Lung

The three organs behind asterixis in an unexplained encephalopathy, worked in the order of the examination: sclerae and abdomen, then when the patient last passed urine and last dialysed, then the respiratory rate and accessory muscles. All three are answered before any investigation returns.

Myoclonus discriminates within that group. Kaufman notes that patients with hepatic or uraemic encephalopathy have asterixis, while myoclonus occurs in uraemia, in penicillin or bismuth intoxication, with meperidine treatment, and in several other metabolic encephalopathies. Uraemia therefore produces both. Tasman widens the differential: myoclonus may be seen in toxic deliria such as the serotonin syndrome, in metabolic deliria such as uraemic encephalopathy and hyperglycaemia, and in certain others such as Hashimoto's encephalopathy. The first arises

from our own prescribing; the last is taught where this chapter deals with subclinical thyroid disease and autoimmunity.

> Asterixis narrows an unexplained encephalopathy to liver, kidney or lung; myoclonus alongside it leans renal.

Discriminating the systemic encephalopathies at the bedside and on the electroencephalogram

CONDITION	ASTERIXIS	MYOCLONUS	ELECTROENCEPHALOGRAM
Uraemic encephalopathy	Yes	Yes	Generalised slowing of the dominant rhythm; triphasic waves when uraemia is severe
Hepatic encephalopathy (taught elsewhere in this chapter)	Yes	Not among those listed	Triphasic waves
Respiratory failure	Yes	Not listed	Not characterised in these sources
Hypoxic-ischaemic encephalopathy	Not listed	Not listed	Generalised periodic discharges

23.4 What the electroencephalogram adds, and the seizure it hides

The electroencephalogram does two jobs. Kaplan and Sadock state that an electroencephalogram can aid diagnosis, typically showing a generalised slowing of the dominant rhythm, and can detect seizures including **nonconvulsive status epilepticus**, which can occur in uraemia and could be mistaken for uraemic encephalopathy. Grading is the first: diffuse slowing confirms the disturbance is global and metabolic rather than focal and structural, the question a confused patient with a normal scan raises. The second changes management, because no bedside observation reliably separates an obtunded, twitching uraemic patient who is encephalopathic from one in nonconvulsive status, and the wrong call leaves a patient in status while dialysis is awaited.

Kaufman records that **triphasic waves** develop in hepatic encephalopathy and in several other toxic-metabolic encephalopathies, such as severe uraemia. The qualifying adjective is load-bearing: Kaufman's worked example shows mild to moderate derangement producing diffuse background slowing rather than triphasic waves, so the waves imply uraemia well beyond the threshold at which the patient first became confused. Their absence excludes nothing.

> **PITFALL**
>
> Accepting a normal or non-specific electroencephalogram report as evidence against a metabolic encephalopathy, when the study was recorded on a lucid day between dialysis sessions. The picture fluctuates and so does the tracing, which cannot answer a question about a state it never sampled.

23.5 Three mimics that live inside the uraemic patient

Uraemic encephalopathy is a diagnosis of exclusion in a population unusually rich in things to exclude. The first mimic is thiamine deficiency. Kaplan and Sadock record that reduced levels of thiamine, a water-soluble vitamin, have been noted in haemodialysis patients, and cases of Wernicke encephalopathy have been mistaken for uraemic encephalopathy. Dialysis is a filter and does not distinguish a neurotoxin from a water-soluble vitamin. A confused, ataxic dialysis patient with eye signs therefore has a treatable deficiency until proved otherwise, and this is not the alcohol-related presentation the syndrome is usually taught through. The acute syndrome and its treatment regimens are taught where this chapter deals with Wernicke encephalopathy and the Korsakoff boundary.

The second is nonconvulsive status, above. The third is radiological. Uraemic encephalopathy has also been associated with a cliniconeuroradiologic syndrome termed **posterior reversible (leuko) encephalopathy syndrome**, whose characteristic findings on computed tomography or magnetic resonance imaging lie in the posterior cortical and subcortical white matter. Imaging here is ordered to exclude a bleed or an infarct and the negative filed as reassurance, yet a positive finding is available in exactly this population: posterior, often subtle on computed tomography, and reversible when recognised.

IN INDIA

Maintenance haemodialysis in much of Indian practice is intermittent by economics rather than by prescription: sessions are missed when money, transport or a working machine is unavailable. Presentations skew late, so overt delirium, seizures and catatonia are commoner first contacts than insidious irritability and insomnia, and the clinician waiting for a gradual prodrome will meet the patient in coma. Electroencephalography is concentrated in teaching and tertiary hospitals, so the investigation that separates encephalopathy from nonconvulsive status is often unavailable. Where it cannot be obtained, hold that question open, escalate dialysis, and do not commit to a psychiatric formulation of a confusional state that has never had a tracing. Thiamine, by contrast, is cheap and everywhere, and empirical use costs almost nothing.

23.6 When dialysis is the cause: disequilibrium and the aluminium dementia

Two encephalopathies arise from the treatment rather than the disease, and candidates confuse them because both are called dialysis something. The first is **dialysis disequilibrium**, an osmotic event. KDT's Essentials of Medical Pharmacology gives its basis as a fall in the osmolality of plasma and extracellular fluid due to rapid haemodialysis or peritoneal dialysis, and lists mannitol as the countermeasure. Goodman and Gilman independently name dialysis disequilibrium syndrome among the indications for mannitol. Beyond that mechanism and its countermeasure these sources give no clean account and no dose; the timing, symptom sequence, severity spectrum and risk factors are stated here as not established rather than filled in from recollection.

The second is the chronic dementing syndrome of aluminium exposure, to which ICD-11 gives a diagnostic home without the traditional name. The code is **6D85.2**, Dementia due to exposure to heavy metals and other toxins. Its essential requirements are that all the diagnostic requirements for dementia are met and that the dementia is presumed attributable to toxic exposure to specific

heavy metals, such as aluminium from dialysis water, lead, mercury or manganese, demonstrated by neuropsychological test data, neuroimaging data, medical tests or clinical history. Two things follow: the classification does not recognise dialysis dementia as an entity, so a candidate searching the manual for that phrase will not find it, and the attribution is presumptive rather than confirmed by any test.

ICD-11 holds that the neurocognitive impairments depend on the particular metal or toxin and can affect any cognitive domain, that onset is related to exposure and progression can be rapid, particularly with acute exposure, and that in some cases symptoms are reversible once the exposure is identified and ceases. A reversible dementia is rare enough to deserve a place on the list whenever a dialysis patient's cognition declines. The dementia syndrome itself and its workup belong to the Dementias and neurocognitive disorders notes.

23.7 Managing uraemic encephalopathy

No guideline-level protocol for the psychiatric management of uraemic encephalopathy appears in these sources; what they give is a treatment logic. **LOCUS.** Once the mental state has changed this is not outpatient work: the patient belongs where dialysis, biochemistry and cardiac monitoring are available, meaning a medical or renal ward with liaison psychiatry attending, and intensive care once there are seizures or a depressed conscious level. The commonest system error is a transfer in the wrong direction, to a psychiatric unit that cannot dialyse, on the strength of a psychotic or catatonic presentation.

FOCUS. Acutely the targets are the derangement and the seizure, not the psychopathology. Kaplan and Sadock give the sequence: removal of uraemic toxins by haemodialysis, correction of electrolyte imbalances and of anaemia, and treatment of malnutrition can diminish symptoms and improve cognition, while seizures may be treated with anticonvulsants. Over the short and medium term the focus shifts to sustaining that correction and to reassessing cognition once the biochemistry settles, because a mental state assessed in the uraemic trough is not the patient's baseline. Long term, definitive renal replacement.

MODUS. Dialysis is the effective modality; anticonvulsants cover the seizure; psychological and social work belongs to adherence and the family burden of sustaining a dialysis schedule, which determines whether the correction holds. Sedation needs the most caution, and the source says why: drugs that are sedating, or that require elimination of the parent compound or an active metabolite by renal clearance, should be avoided, or their dosages appropriately adjusted for the degree of impaired renal function. Making that adjustment agent by agent is taught where this chapter deals with prescribing psychotropics in hepatic and renal impairment.

Definitive treatment does not guarantee a clean cognitive result. Kaplan and Sadock report a cross-sectional study of kidney transplant recipients an average of 3.4 years afterwards in which **58%** met criteria for cognitive impairment. That figure describes transplant recipients at that interval and nothing else; it is not a prevalence in uraemia, and quoting it as one is an error of population.

58%	3.4 YEARS	3	6D85.2
COGNITIVELY IMPAIRED AFTER TRANSPLANT	MEAN INTERVAL IN THAT STUDY	METABOLIC DELIRIA WITH ASTERIXIS	ICD-11 CODE, DIALYSIS ALUMINIUM

TRAP **Sending a serum ammonia to diagnose uraemic encephalopathy.** Two standard sources disagree and the disagreement is unresolved; know both positions. Kaplan and Sadock name the uraemic neurotoxins, insist no single metabolite is the sole cause, and never invoke ammonia in the renal context. Against that, a widely reproduced differential table in Saunders and colleagues' Addiction medicine lists serum ammonia among the differentiating features of uraemic encephalopathy. That table shows clear row displacement, one value repeated verbatim down five unrelated rows and another carrying two competing values, so the ammonia entry may have migrated from the adjacent hepatic line. Ammonia is the hepatic marker, its role in uraemia is unsupported by the major textbook account, and no metabolite test diagnoses this condition.

23.8 Hypoxia and hypercapnia: the encephalopathy referred as anxiety

Respiratory failure is the third organ behind asterixis, and the one most often handed to psychiatry as a psychiatric problem. Kaufman puts it as a vignette: a man with rapidly advancing Guillain-Barre syndrome develops confusion, overwhelming anxiety and agitation, and the answer is that he may be developing hypoxia, hypercapnia or both because of chest and diaphragm muscle paresis. Guillain-Barre syndrome, also called acute inflammatory demyelinating polyradiculoneuropathy, is not directly associated with central nervous system dysfunction; however, respiratory insufficiency is a common complication of the illness and might cause the anxiety and agitation.

That is hypoxic and hypercapnic encephalopathy arriving under a psychiatric label. Acute anxiety and agitation in any patient with neuromuscular or respiratory compromise is a blood-gas problem until proved otherwise, and the request that arrives on the phone is almost always for a benzodiazepine, which sedates the respiratory drive that is the only thing keeping the gases survivable. The first response is an arterial blood gas and an assessment of ventilation; the second, if the gases are deranged, a call to intensive care rather than a prescription. ICD-11's list of conditions to consider in a disturbed mental state includes Wernicke's encephalopathy, electrolyte disturbance, systemic infection, and hypoxia or hypercapnia: a bare list rather than a diagnostic rule, quoted only to place the gases in the standard differential.

GOOD TO REMEMBER

Arterial carbon dioxide is a less reliable guide to encephalopathy than arterial oxygen, because chronic retainers adapt and tolerate levels that would confuse an acutely deranged patient. The same principle governs uraemia: an acute derangement produces symptoms at a degree of abnormality a slowly progressive one is tolerated at. A single gas read without the trajectory decides nothing. The thresholds conventionally quoted for the two partial pressures could not be verified from a clean reading of these sources and are deliberately not printed.

23.9 Anoxic injury, its end states, and the encephalopathy treatment causes

Anoxia is the exception to the reassuring rule that metabolic encephalopathies reverse: most patients recover from delirium if the underlying illness responds to treatment, but sometimes, as

in anoxic encephalopathy, they are left with permanent cognitive impairment. To a neurologist **watershed** denotes areas of cerebral cortex perfused by the terminal branches of arteries, whose already tenuous blood supply falls to an insufficient level during hypotension or anoxia. What remains once the insult ends is structural loss, not biochemical disturbance.

Kaufman describes a severe and extensive, though incomplete, cerebral cortical anoxic insult, an injury also suffered by survivors of cardiac arrest and by those who have attempted suicide by hanging or carbon monoxide poisoning, in whom the well-perfused language arc may be preserved while almost the entirety of the other cortical regions is destroyed. That preserved arc surrounded by destroyed cortex is the **isolation aphasia** picture: the patient who repeats fluently but comprehends nothing and initiates nothing. Two of the three causes named are suicide attempts, which is why the survivors are our patients.

Where the cortical injury is more complete the outcome is a **persistent vegetative state**, which typically follows cerebral cortical anoxia from cardiac arrest, drug overdose, carbon monoxide poisoning or traumatic brain injury. The cortical injury is extensive, so cognitive impairment is profound, yet vegetative functions such as respiration and digestion continue and the eyes open and close, mostly randomly or in response to light. That extensive cortical damage produces slow and disorganised electroencephalographic activity. That tracing discriminates against locked-in syndrome, which follows infarction at the base of the lower brainstem or extensive cranial and peripheral nerve disruption: hemispheres and thalamus intact, those patients retain normal cortical function and a normal tracing. At the far end, electroencephalographic activity is completely absent, or **electrocerebral silence**, in brain death, in which by definition brain function is utterly abolished.

The surviving anoxic patient often shows a periodic pattern. **Generalised periodic discharges** are commonly seen in toxic or metabolic encephalopathies including sepsis, where triphasic waves are especially characteristic; in extensive structural cerebral damage such as hypoxic ischaemic encephalopathy, diffuse intracranial haemorrhage or neurodegeneration; in the epileptic encephalopathies; with medications; and in infections such as Creutzfeldt-Jakob disease or subacute sclerosing panencephalitis. The medication list is the one to memorise, because one of its members is ours:

- Cefepime.
- Lithium.
- Ifosfamide.
- Baclofen.
- Metrizamide.

Lithium as a drug, its pharmacokinetics, therapeutic range and non-thyroid organ-system adverse effects belong to the Mood disorders: pharmacotherapy notes.

TRAP **Reading generalised periodic discharges as an independent marker of death.** Candidates reproduce the traditional teaching that they imply a poor prognosis. The source describing the pattern corrects it: recent studies have not found an independent association with increased mortality; the discharges are associated with nonconvulsive status epilepticus, which does itself carry increased mortality; and the underlying aetiology and the severity of the primary disease process are most often the key determinants of prognosis. The pattern is a prompt to look for treatable status and to characterise the primary illness, not a prognostic verdict to deliver to a family.

One systemic encephalopathy is produced entirely by treatment. Kaufman warns that overly rapid correction of metabolic derangements, such as hyponatraemia, can lead to **central pontine myelinolysis**, the prime example of the osmotic demyelination syndrome. The syndrome itself, the populations at risk, the safe correction rate and the management of the hyponatraemia that provokes it belong where this chapter deals with psychogenic polydipsia and the sodium-correction risk; no numeric ceiling for the rate is stated in these sources and none is printed here. The principle: a chronic derangement must be corrected slowly, because it is the speed of correction and not the abnormality that demyelinates.

24 · Prescribing psychotropics in hepatic and renal impairment

This is the commonest real prescribing problem in liaison psychiatry and one of the easiest places to lose marks. Most answers turn on one insight: the drugs safest in liver disease are the drugs most dangerous in kidney disease.

Two standard references dominate and do not always agree; where they diverge, both positions are given rather than harmonised. Only the organ-impairment modifier is owned here: dose and titration for the psychiatric indication belong with that disorder's pharmacotherapy, lithium as a drug to the Mood disorders: pharmacotherapy notes, and antipsychotic selection to the Schizophrenia: management, antipsychotics and adverse effects notes.

ONE-THIRD	4%	0.5% TO 3%	1.65%
ON ANTIPSYCHOTICS, ONE ABNORMAL LFT	AN LFT OVER THREE TIMES ULN	ON ANTIDEPRESSANTS, TRANSAMINASE RISE	SEVERE HEPATIC TOXICITY, QUETIAPINE

24.1 What the liver function tests do not tell you

Liver function tests evaluate hepatic damage; they say little about **hepatic metabolising capacity**, and many people with chronic liver disease are asymptomatic or fluctuate clinically. A normal panel does not license a normal dose, nor a deranged one a reduced dose. There are few clinical studies of psychotropics in hepatic disease, so almost all of what follows is kinetic inference rather than trial evidence.

Advanced organ disease shifts perfusion, fluid status, protein binding, hepatic metabolism and excretion, with pretransplant patients additionally prone to encephalopathy, arrhythmia and drug interactions. Two dosing consequences follow. Highly protein-bound drugs, the tricyclics, the SSRIs other than citalopram, trazodone and the antipsychotics, may carry raised free levels initially, and this is not reflected in the measured total plasma level, which is why therapeutic

drug monitoring misleads here. Second, agents with extensive **first-pass metabolism**, the tricyclics and haloperidol, need lower doses.

24.2 The principles, and the four drugs that escape the liver

The general principles of prescribing psychotropics in hepatic impairment generate most correct answers without a table.

- Be cautious with extensively hepatically metabolised drugs, which is most psychotropics; lower doses may be needed.
- The exceptions are sulpiride, amisulpride, lithium and gabapentin, undergoing no or minimal hepatic metabolism.
- Avoid very sedative drugs, for the risk of precipitating hepatic encephalopathy.
- Avoid very constipating drugs, for the same reason.
- Avoid drugs hepatotoxic in their own right, the examples given being the monoamine oxidase inhibitors and chlorpromazine.

Both rules earn their place: sedation and constipation slow clearance of ammoniagenic load from the gut and mask worsening encephalopathy. That syndrome is set out where this chapter deals with hepatic encephalopathy, and delirium belongs to the Stroke, TBI, delirium and focal brain syndromes notes. Clozapine is disqualified twice over, being both.

RULE **Reach for the drug the failing organ does not handle.** In hepatic impairment the answers are the renally cleared agents: amisulpride and sulpiride need no reduction provided renal function is normal, and paliperidone, mainly excreted unchanged, may be a good choice in pre-existing hepatic disease. Turn the failing organ around and the list inverts: be cautious with extensively renally cleared drugs such as sulpiride, amisulpride and lithium, and avoid nephrotoxic drugs, lithium and the non-steroidal anti-inflammatories among them.

The four drugs that escape the liver are sulpiride, amisulpride, lithium and gabapentin. They are the hepatic-impairment answers and, gabapentin partly excepted, the renal-impairment dangers. Learn the list once and read it both ways.

24.3 How often psychotropics disturb the liver tests

Abnormal transaminases on a psychotropic are common; significant injury is not. On antipsychotics, a third of patients have at least one abnormal test and 4% have one above three times the upper limit of normal, transaminases most often affected, generally within 1 to 6 weeks of initiation. On antidepressants, 0.5% to 3% develop asymptomatic mild transaminase elevation, normally several days to 6 months after initiation, the elderly more vulnerable. Frank damage is rare and mostly **idiosyncratic**, being unpredictable, unrelated to dose, and with cross-toxicity within a class described. Within-class substitution after a hepatic reaction is therefore not automatically safe.

Transient asymptomatic rises appear in 20% of patients on haloperidol, and on clozapine elevations of AST, ALT and GGT to over twice normal occur in up to a third, resolving

spontaneously in 6 to 12 weeks. Only one agent carries a mandated schedule: agomelatine needs liver function tests at baseline and at 3, 6, 12 and 24 weeks during initiation and at each dose increase, stopping if transaminases exceed three times the upper limit of normal.

TRAP **The consequence rises, not the probability.** Pre-existing liver disease does not increase the risk of drug-induced hepatotoxicity, but it may be more catastrophic if it occurs. Candidates reverse this and answer that cirrhosis makes hepatotoxic reactions likelier. The corollary: shun a drug with a hepatic signal in the patient with no reserve to absorb the hit.

24.4 Antidepressants in hepatic impairment

Nearly every antidepressant is hepatically cleared, so most differences are of degree. Three are outright prohibitions. Duloxetine's clearance is markedly reduced even in mild disease, with hepatocellular injury beyond ten times the upper limit of normal and fatal hepatic failure reported; agomelatine likewise; non-selective monoamine oxidase inhibitors are contraindicated in hepatic impairment.

AGENT	INDICATION	MODIFIER IN HEPATIC IMPAIRMENT
Sertraline	Depression, anxiety	Low or less frequent dose in mild disease; avoid in moderate and severe. Itself used for **cholestatic pruritus**
Citalopram	Depression	Maximum **20 mg daily**; higher exposure, greater QT risk
Escitalopram	Depression	60% higher exposure in mild to moderate disease; 5 mg daily for the first 2 weeks, maximum 10 mg daily
Fluoxetine	Depression	Accumulates in compensated cirrhosis; reduce by at least 50% or dose on alternate days
Mirtazapine	Depression	50% dose reduction on kinetic data, though sedative
Agomelatine	Depression	Contraindicated, including cirrhosis and active liver disease
Moclobemide, reboxetine	Depression	Moclobemide to half or one-third of the dose; reboxetine a 50% reduction in starting dose
Venlafaxine	Depression	Reduce the total daily dose; watch blood pressure
Bupropion	Depression	In severe disease, 150 mg extended-release every other day or 100 mg per day sustained-release

AGENT	INDICATION	MODIFIER IN HEPATIC IMPAIRMENT
Vortioxetine, esketamine	Depression	Neither recommended in severe disease; esketamine needs longer monitoring in moderate disease

Sertraline is where the references part company: one permits reduced or less frequent dosing at any severity, the other prohibits it in moderate and severe disease. Sertraline being otherwise a natural first choice in the medically ill, the disagreement is real and unresolved, compounded because the severity scores printed against it overlap and cannot both be correct.

24.5 Antipsychotics in hepatic impairment

Aripiprazole needs no reduction in mild to moderate disease, with caution in severe disease, and pimavanserin none at all. Clozapine is contraindicated in active liver disease with nausea, anorexia or jaundice, in progressive disease and in hepatic failure; in milder disease start treatment-resistant schizophrenia at 12.5 mg, increasing slowly, plasma levels gauging metabolising capacity. Olanzapine changes little kinetically even in severe impairment but is sedative and anticholinergic, so start psychosis or mania at 5 mg per day in moderate or severe disease, plasma target 20 to 40 mcg per litre; with clozapine it is the antipsychotic most often linked to drug-induced liver injury. For psychosis in severe disease, risperidone starts at 0.5 mg twice daily, rising by 0.5 mg twice daily no faster than weekly above 1.5 mg twice daily.

In moderate or severe hepatic failure, brexpiprazole should not exceed 3 mg per day for schizophrenia or 2 mg per day for depression or agitation in Alzheimer's disease, its half-life long at around 90 hours. For schizophrenia or bipolar depression, lurasidone starts at 18.5 mg, capped at 74 mg per day in moderate disease, where exposure rises 1.7-fold, and 37 mg per day in severe disease, where it trebles; the American reference gives the same ceilings in salt-equivalent strengths, starting at 20 mg per day and not exceeding 80 mg per day in moderate or 40 mg per day in severe hepatic impairment. For psychosis, briefly: asenapine is avoided in severe disease, exposure rising sevenfold; iloperidone is reduced in moderate disease, active metabolites doubling, and avoided in severe; cariprazine is not advised in severe failure, half-life about 2 to 4 days; lumateperone has a manufacturer-recommended dose of 21 mg daily in moderate and severe disease, exposure being increased; and the phenothiazines, sedative and constipating to a drug, are best avoided, chlorpromazine being particularly hepatotoxic.

TRAP **Haloperidol and quetiapine are where the two standard references give opposite advice.** For psychosis, agitation or delirium, one position holds haloperidol extensively hepatically metabolised, so halve initial doses and adjust in smaller increments at longer intervals; the other records that typical agents such as haloperidol need no dosage adjustment in renal or hepatic impairment. Each is internally consistent, one reasoning from first-pass extraction, the other from the near-absence of renal excretion. Quetiapine splits identically: one holds that clearance falls by a mean of 30%, so psychosis or bipolar disorder starts at 25 mg per day of the immediate-release or 50 mg per day of the extended-release preparation, the other that no adjustment is needed. The sources also print two different lurasidone starting figures, and impairment tables have carried a millilitre where a milligram was meant; an oral tablet dose is never in millilitres.

24.6 Sedatives, mood stabilisers, transplantation and the substance-use agents

Benzodiazepine choice in liver disease follows metabolic route, not half-life. Impaired hepatic oxidation delays clearance and raises adverse effects, so agents needing only **conjugation** are preferred for anxiolysis or sedation in hepatic impairment: lorazepam, oxazepam and temazepam, with respiratory depression the residual risk.

MNEMONIC

Lorazepam, Oxazepam, Temazepam

LOT: the three benzodiazepines needing no hepatic oxidation, and so the choices in liver disease. In the cirrhotic patient in withdrawal, offered diazepam or chlordiazepoxide alongside these three, this is the answer.

Among anticonvulsant mood stabilisers the most examinable adjustment: for bipolar disorder or epilepsy in severe hepatic impairment, lamotrigine has initial, escalation and maintenance doses reduced by **50%** with ascites and by 25% without. Valproate is the conspicuous hole: both references say only to avoid it where possible, since clearance falls and low albumin raises the free fraction, and neither offers a percentage or a ceiling.

Transplantation adds an interaction problem to the kinetic one. Immunosuppressants are largely CYP3A4 substrates and psychotropics differ widely in inhibiting it: in decreasing order, fluvoxamine and nefazodone, then fluoxetine, then sertraline, the tricyclics and paroxetine, then venlafaxine. Nefazodone has produced toxic immunosuppressant levels and is best avoided outright. Induction is no safer: carbamazepine and oxcarbazepine, as CYP3A4 inducers, lower immunosuppressant levels and risk organ rejection.

GOOD TO REMEMBER

Before starting any antidepressant in a transplant candidate or recipient, check the immunosuppressant regimen and speak to the transplant team. The commonest avoidable harm is not hepatotoxicity but a graft lost to a predictable interaction.

Naltrexone carries a black box warning for hepatotoxicity and is not recommended in active liver disease, hepatitis or liver failure; above the recommended range, at more than 300 mg per day, it

becomes a direct hepatotoxin. Methadone needs a lower initial dose and slow titration in either impairment and can prolong the QT interval. For ADHD, atomoxetine is reduced by 50% in moderate and 75% in severe hepatic impairment, modafinil by 50% in severe, while methylphenidate and amphetamine need no adjustment in either organ.

24.7 Estimating renal function, and the general renal principles

CKD-EPI is more accurate than MDRD and is preferred; for most drugs and most adults of average build its estimated glomerular filtration rate can guide dose adjustment. For nephrotoxic drugs, patients aged 75 years and over, and either extreme of muscle mass, body mass index below 18 or above 40 kg/m2, calculate **creatinine clearance** instead, by the **Cockcroft and Gault equation**. Regulatory advice extends creatinine clearance to direct-acting oral anticoagulants and to narrow therapeutic index drugs mainly renally excreted, lithium among them. The psychiatric point: the drug most likely to need renal dosing is the one for which the routinely reported estimate is the wrong number.

Five principles frame dose adjustment in renal impairment and the choice of agent. renal function significantly affects elimination when the fraction excreted unchanged in urine is **30% or more** of the dose, which is why lithium, sulpiride, amisulpride and paliperidone dominate the renal tables and clozapine and haloperidol barely appear. Age is a silent modifier: adults over 65 should be assumed to have at least mild renal impairment, serum creatinine perhaps not raised because muscle mass is smaller. Renal failure slows hepatic clearance too, the point most often missed: elimination of hepatically metabolised drugs can fall in kidney disease, possibly through inhibition of enzyme activity by **uraemia**, so start low and go slowly, half-life and time to steady state often being prolonged. Long-acting preparations such as depots should be avoided, since their dose and frequency cannot easily be adjusted if renal function changes, as should QTc-prolonging drugs, electrolyte disturbance being common in established renal failure. Trajectory matters as much as the current value: a 30% change over 2 years carries a fivefold increase in the risk of end stage renal disease.

> **PITFALL**
>
> Manufacturer bands are written variously in filtration rate or in creatinine clearance, and for a narrow therapeutic index drug the two are not interchangeable. The pragmatic bands used by the drug tables also do not map cleanly onto the chronic kidney disease stages, so do not translate between them.

24.8 Lithium, and the drugs the kidney clears

Lithium is the paradigm case, being both the drug most affected by renal impairment and a cause of it. It is nephrotoxic and contraindicated in severe renal impairment, **95%** excreted unchanged in urine. For bipolar maintenance at a filtration rate of 10 to 50 mL/min, avoid it or reduce to 50% to 75% of the normal dose with level monitoring; below 10 mL/min avoid if possible, and if used, 25% to 50% is essential. It can be used successfully during haemodialysis. The renal risk is cumulative rather than acute: risk factors include increasing age, treatment duration, cumulative dose, lower initial eGFR, female sex, hypertension and diabetes, concomitant nephrotoxic drugs,

nephrogenic diabetes insipidus and previous toxicity, and damage is likelier with chronic than acute toxicity. Prevention rests on once daily dosing, tight adherence to recommended plasma levels, avoiding intoxication, treating comorbidity assertively and monitoring kidney function actively.

IN INDIA

Wherever lithium is prescribed, the creatinine clearance instruction becomes a daily problem: laboratories report an estimated filtration rate with every creatinine, while Cockcroft and Gault needs a recorded body weight that outpatient notes frequently lack. Non-steroidal anti-inflammatories, among the nephrotoxic drugs to avoid where renal reserve is limited, are bought across the counter and rarely volunteered unless asked for by name. And long-acting injectables, attractive where follow-up is distant and adherence fragile, are the preparations to avoid when renal function may change, because the dose cannot be taken back.

AGENT	INDICATION	MODIFIER BY RENAL FUNCTION
Amisulpride	Psychosis	50% unchanged. 50% of the dose at GFR 30 to 60 and 33% at 10 to 30 mL/min; no recommendation below 10, so avoid
Sulpiride	Psychosis	Almost wholly renal. 70%, 50% and 34% of the normal dose at GFR 30 to 60, 10 to 30 and below 10 mL/min, or lengthen the interval 1.5-, 2- and 3-fold
Paliperidone	Psychosis	59% unchanged. GFR 50 to 80: 3 mg daily, maximum 6 mg. GFR 10 to 50: 3 mg alternate days or 1.5 mg daily. Clearance falls 71% in severe disease; oral contraindicated below GFR 10, depots below 50 mL/min
Risperidone	Psychosis	Clearance falls 60%. Below GFR 50: 0.5 mg twice daily for at least 1 week, then up by 0.5 mg twice daily to 1 to 2 mg twice daily; injection only after oral titration
Olanzapine	Psychosis	7% unchanged. UK labelling: 5 mg daily below GFR 50, injection from 150 mg every 4 weeks. US labelling: no adjustment
Lurasidone	Schizophrenia, bipolar depression	9% unchanged; levels 1.5-, 1.9- and 2.0-fold higher in mild, moderate and severe disease. Below GFR 50: start 18.75 mg, maximum 74 mg per day; avoid below 15 mL/min

AGENT	INDICATION	MODIFIER BY RENAL FUNCTION
Clozapine, chlorpromazine, aripiprazole, lumateperone, pimozide, haloperidol, quetiapine	Psychosis, delirium	All excreted overwhelmingly as metabolites, under 1% of haloperidol and under 5% of quetiapine unchanged, and none needs routine adjustment. Clozapine and chlorpromazine: no adjustment above GFR 10 mL/min, cautious titration in very severe disease, though clozapine is contraindicated by its manufacturer in severe renal disease. Haloperidol: normal dosing at GFR 10 to 50, lower start below 10. Quetiapine clearance falls 25% below GFR 30 with no change advised
Ziprasidone	Psychosis	Under 1% unchanged, no adjustment above GFR 10, but the injection carries **cyclodextrin sodium**, a renally eliminated excipient
Valproate	Bipolar disorder, epilepsy	About 2% unchanged; normal dosing, but free levels rise and below GFR 10 dose to free valproate levels. Interstitial nephritis and renal tubular acidosis reported
Topiramate, gabapentin, oxcarbazepine	Bipolar disorder, epilepsy, neuropathic pain	Topiramate halved in moderate to severe disease. Gabapentin 600 mg per day in moderate and 150 to 300 mg per day in severe disease, supplemented after dialysis. Oxcarbazepine from 300 mg per day in severe disease
Desvenlafaxine, acamprosate, varenicline, buspirone	Depression, dependence, anxiety	Desvenlafaxine no more than 50 mg per day in moderate and 50 mg every other day in severe disease. Acamprosate 333 mg three times daily in moderate disease, contraindicated in severe; varenicline from 0.5 mg per day to 0.5 mg twice daily in severe disease. Buspirone not recommended in severe renal or hepatic impairment

Two points close the table. Amisulpride and sulpiride, always taught together, carry different percentage schedules at the same bands and must be learned apart. And risperidone warns about triggers, not doses: three references give the same 0.5 mg twice daily start for psychosis but

attach it to a filtration rate below 50 mL/min, to severe impairment, and to a creatinine clearance below 30 mL/min respectively, materially different populations.

24.9 Choosing a regimen and following it up

Management here is a sequence, not a lookup. The LOCUS of care is set by the organ, not the psychiatric diagnosis: decompensated cirrhosis or established renal failure needing psychotropic initiation belongs on a medical or liaison inpatient service where electrolytes and fluid balance are already followed; stable compensated disease sits in outpatient psychiatry under a shared plan with a named physician; agitation in the encephalopathic patient is handled where the medical team is. Three bodies of guidance frame the decision: the Maudsley Prescribing Guidelines for drug-by-drug kinetic adjustment, Kaplan and Sadock's Comprehensive Textbook of Psychiatry for the transplant and advanced organ disease perspective, and the American Psychiatric Association practice guideline where it specifies initial doses and titration rates in impairment.

The FOCUS moves by phase: acutely, the symptom threatening the medical course, at the smallest effective exposure of the safest molecule; in the short term, whether the syndrome is primary or a manifestation of organ failure, which decides whether the drug continues at all; in the mid term, titration against a moving kinetic target, time to steady state being prolonged; in the long term, surveillance of the organ itself.

The MODUS. Pharmacologically, choose by clearance route first and receptor profile second, and lengthen the dosing interval rather than shaving a tablet where licensed strengths do not divide. Psychological treatment carries no hepatic or renal penalty and is underused in exactly these patients; electroconvulsive therapy remains available when the pharmacological field has narrowed to nothing safe; socially, address alcohol, over-the-counter analgesia and dehydration. One instruction generalises: pretransplant patients call for lower starting dosages and slower titration until clinical effectiveness is reached.

25 · Reproductive hormones and mood: PMDD, perimenopause and the pituitary

Four clinical territories sit under one axis. Premenstrual dysphoric disorder, perimenopausal depression, polycystic ovary syndrome with its prolactin and pituitary differential, and androgen-related mood disturbance share one hypothalamic-pituitary-gonadal axis. The marks sit at the thresholds, where the sources contradict each other and sometimes themselves.

1.3%	2 TO 5-FOLD	10.5%	5%
PMDD BY PROSPECTIVE DAILY RATING, UNITED STATES	RISE IN DEPRESSION RISK DURING THE PERIMENOPAUSE	INCIDENT PCOS ON VALPROIC ACID VERSUS 1.4% ON OTHER MOOD STABILISERS	MANIC OR HYPOMANIC SYNDROMES UNDER BLINDED HIGH-DOSE ANDROGEN EXPOSURE

25.1 The axis, and why change matters more than level

The hypothalamus releases gonadotropin-releasing hormone in a pulsatile fashion, driving luteinising hormone and follicle-stimulating hormone from the anterior pituitary and then gonadal oestrogen, progesterone and testosterone. Oestrogen rises through the follicular phase

until a critical concentration triggers the luteinising hormone surge, and progesterone rises in the luteal phase; fluctuating oestrogen is held to dysregulate monoamines, producing luteal-phase mood and anxiety symptoms. The axis belongs to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

Progesterone becomes **allopregnanolone**, and serotonergic antidepressants are believed to enhance gamma-aminobutyric acid type A receptor sensitivity, or promote allopregnanolone formation by increasing 3-alpha-hydroxysteroid dehydrogenase activity. Testosterone runs the same route, its metabolite **3-alpha-androstanediol** likewise a positive allosteric modulator of gamma-aminobutyric acid receptors.

RULE **Mood follows the change in hormone concentration and the individual's sensitivity to it, not the absolute level.** It is the change in hormone levels rather than the levels themselves that triggers the symptomatic state. Under blinded conditions, acute oestradiol withdrawal induced depressive symptoms in asymptomatic postmenopausal women with a past perimenopausal depression, but not in those without. Same hormone, opposite outcome, decided by who the woman is.

25.2 Premenstrual dysphoric disorder: the criteria and their arithmetic

Criterion A fixes timing and count: in the majority of cycles, **at least five symptoms** must sit in the final week before menses, improve within a few days after menses begin, and become minimal or absent in the week postmenses. Criterion B makes an affective symptom compulsory: one or more of marked affective lability, marked irritability or anger with increased interpersonal conflict, marked depressed mood with hopelessness, or marked anxiety and tension.

Criterion C supplies the remainder: decreased interest in usual activities, difficulty concentrating, lethargy, marked appetite change with food cravings, hypersomnia or insomnia, feeling overwhelmed, and physical symptoms such as breast tenderness, joint or muscle pain, bloating or weight gain. These count together with the Criterion B symptoms to reach a total of five, never five from Criterion C alone.

MNEMONIC

Lability, Irritability, Depressed mood, Anxiety

The four Criterion B items, at least one mandatory. Somatic symptoms cannot carry it: bloating, breast tenderness, hypersomnia, food craving and fatigue is five symptoms and no diagnosis.

Criterion D requires clinically significant distress or interference with work, school, usual social activities or relationships; Criterion E excludes what is merely an exacerbation of another disorder, though co-occurrence is allowed; Criterion G excludes a substance, a medication or another medical condition, with hyperthyroidism the worked example, the hook to the hypothyroidism and hyperthyroidism in psychiatry notes. A note requires that Criteria A to C have been met for most cycles in the preceding year.

25.3 Confirming it, and the three things it is not

Criterion F makes this a measurement diagnosis: Criterion A should be confirmed by prospective daily ratings during at least two symptomatic cycles, though the diagnosis may be made provisionally beforehand, with the word provisional appended. The validated prospective instrument is the **Daily Rating of Severity of Problems**; the Premenstrual Symptoms Screening Tool is retrospective and screening only.

IN INDIA

The two-cycle rating requirement fails here structurally, not clinically: the woman with cyclical irritability is seen once and will not return two months later with a completed diary, and literacy and privacy compound it. Use the manual's own provision, since a provisional diagnosis before rating confirmation is explicitly permitted, and treat the retrospectively-screened patient as a candidate for rating, not a confirmed case.

CONDITION	WHAT DEFINES IT	THE DISCRIMINATING OBSERVATION
Premenstrual syndrome	Requires neither a minimum of five symptoms nor mood symptomatology, and is generally less severe	Prevalence about 20% in the manual against 80 to 95% in a text counting any premenstrual symptom. Different constructs, so neither is the prevalence.
Premenstrual exacerbation	Luteal worsening of a pre-existing mood or anxiety disorder, in approximately 50 to 66% of women	After menses start, residual depression and anxiety symptoms persist. No symptom-free postmenstrual interval rules the diagnosis out.
Premenstrual dysphoric disorder	The full criteria set, confirmed prospectively	Symptoms confined to the luteal phase, resolving completely when the period starts.

25.4 How common, and what each number counts

A large German study gives a 12-month community prevalence of 5.8%; a United States study following two cycles found 1.3%, the most rigorous estimate, requiring impairment and no comorbidity, since retrospective estimates are often higher than those based on prospective daily ratings. The wider published band is 1.3 to 7.7%. Heritability is a grammatical trap: no studies have examined heritability in the disorder itself, and estimates of **30% to 80%** apply to premenstrual dysphoric symptoms, where it remains unclear whether the symptoms are heritable or simply associated with other heritable traits. Symptoms cease after the menopause, though cyclical hormone replacement can trigger their re-expression.

25.5 Treating premenstrual dysphoric disorder

LOCUS. Outpatient and shared, since the contraceptive option here is usually prescribed by primary providers and gynaecologists. The Royal College of Obstetricians and Gynaecologists guideline places selective serotonin reuptake inhibitors first line, supported by a meta-analysis finding similar efficacy between agents, and also places the combined oral contraceptive pill first

line, prescribed cyclically or back to back, typically for 3 months. Three antidepressants are approved by the United States Food and Drug Administration, all selective serotonin reuptake inhibitors - fluoxetine, controlled-release paroxetine and sertraline - plus the combined oral contraceptive drospirenone with ethinyl oestradiol, though all agents in the class could be used and venlafaxine extended-release has also been shown effective. Every approval named in this topic is a United States one; Indian licensing status is a separate question and none of these sources state it.

FOCUS. Response is fast and partial: improvement may be seen during the first week, and mood, irritability and anxiety improve while physical symptoms such as swelling, bloating and breast tenderness are less responsive. Warn the patient, or a working drug reads as failure.

MODUS. Starting doses are all the sources supply. Sertraline is usually started at 50 mg per day, with upward titration not uncommonly required. Fluoxetine is often started at 10 to 20 mg per day. Controlled-release paroxetine starts at 12.5 mg per day, as a single morning dose, daily or only during the luteal phase. No source states a target or ceiling, so titrate to response. Intermittent treatment for the last 2 weeks of the cycle can be as efficacious as continuous administration and does not produce discontinuation symptoms; call that window roughly the two weeks from ovulation to menses, since no source defines it by cycle number.

The hormonal ladder: the approved contraceptive is dosed daily for 24 days, followed by 4 days of placebo, but a 2021 meta-analysis reported that combined oral contraceptives improved overall premenstrual symptomatology but not premenstrual depressive symptoms. Second line is continuous percutaneous bioidentical oestradiol, with a progestogen in women with a uterus; beyond it, **gonadotrophin-releasing hormone agonists** give a reversible medical menopause, used therapeutically and as a diagnostic test. Suppression works in about two-thirds of women, as many as are successfully treated by selective serotonin reuptake inhibitors though not necessarily the same individuals, while non-response in about 30% despite adequate ovarian suppression makes a reproductive hormone aetiology unlikely, with a ceiling of not greater than 3 months without gynaecology consultation. Adjunctively, cognitive behavioural therapy is effective, anxiety or irritability may warrant as-needed benzodiazepines, buspirone or hydroxyzine, and zuranolone is being studied rather than approved here. There is no allopregnanolone assay to order.

25.6 Staging the menopausal transition, and the follicle-stimulating hormone question

The menopause is a date in time, 12 months after the last menstrual period, normally between 45 and 55 years; the perimenopause is the phase before it, with erratic ovarian function, symptoms typically arising 2 to 7 years beforehand. Two variants matter: early menopause before 45 years and premature ovarian insufficiency below 40 years, against the instruction that menopausal symptoms should be considered in women older than 40. The transition lasts a median of 4 years, with median age of onset 47 and the final menstrual period at 51 to 52. Staging is menstrual: the 2012 revision of the **Stages of Reproductive Aging Workshop** criteria define the early transition by persistent differences of 7 days or more in the length of consecutive cycles and the late transition by intervals of amenorrhoea longer than 60 days.

TRAP **There is no free-standing follicle-stimulating hormone cut-off for the menopause, and a question offering a bare number is testing whether you know that.** The staging criteria treat elevated plasma follicle-stimulating hormone above 25 IU/L as confirmatory evidence only. The prescribing guidelines answer a different question: in women younger than 45, two raised levels above 30 mIU/mL taken 6 weeks apart, sampled between the second and fifth days of the cycle, may help confirm a diagnosis if clinical symptoms are not conclusive, the Greene Climacteric questionnaire being sufficient in most women older than 40. A psychopharmacology text rejects measurement outright, since hormone testing is neither necessary nor recommended, because these hormones fluctuate during the same day. The two units are numerically the same, so the thresholds genuinely differ. Construct and population separate them: one confirms perimenopause, one confirms menopause in women younger than 45 only, one denies single measurements any value.

25.7 Perimenopausal depression and what to prescribe

There is a 2 to 5-fold increase in the risk of depression during the perimenopause, and suicidal ideation is present in around 6% of women experiencing menopause-related psychological symptoms. Lifetime prevalence is 28% in women with no previous history of major depressive disorder and 59% in women with such a history. Nosologically, research strongly suggests perimenopausal depression is an entity of its own, but the manual contains no specifier or classification for it, so the clinician must use the diagnosis of major depressive disorder.

TRAP **"Is hormone replacement a treatment for perimenopausal depression?" has three published answers, two from the same textbook, and quoting any one alone is the error.** The prescribing guidelines call hormone replacement first-line for menopausal insomnia, anxiety and depression. The comprehensive textbook, on antidepressant augmentation, reports that recent studies of oestrogen augmentation of selective serotonin reuptake inhibitors in both pre- and postmenopausal women have not shown a consistent benefit. The same book, on gonadal steroids, reports that several randomised placebo-controlled trials of oestradiol therapy in perimenopausal depression have demonstrated both short-term antidepressant efficacy and preventative benefits. Population and role separate them: oestrogen as a general antidepressant, or as augmentation in depressed women at large, fails; oestradiol in symptomatic perimenopausal women has trial support.

The usable version is narrow. Oestrogen with the progestin medroxyprogesterone acetate may increase coronary heart disease and stroke risk during the first year, so estradiol prophylaxis can no longer be recommended, but it may still be indicated in women with perimenopausal symptoms and depression who fail standard antidepressant treatment. The preparation suggested is the 50 microgram transdermal estradiol patch, giving levels slightly above the 60 pg/mL threshold necessary to preserve bone: a patch strength and a serum concentration, never merged. In a woman with a uterus endometrial protection is not optional, since unopposed oestrogen for more than 1 year gives a 2 to 12-fold increased relative risk of endometrial cancer, the optimal regimen being 12 to 13 days of progestin treatment each month when oestrogens are administered. On antidepressants, the only antidepressant approved by the United States Food and Drug Administration for perimenopausal vasomotor symptoms is paroxetine, with off-label venlafaxine and desvenlafaxine, and also oestrogen replacement and a neurokinin 3 receptor antagonist, while perimenopausal depression itself can be treated with all approved

antidepressants for major depressive disorder. Approved antidepressant, not approved treatment: that is the trap. Antipsychotic selection in these women belongs to the Schizophrenia: management, antipsychotics and adverse effects notes.

25.8 Polycystic ovary syndrome and the psychiatric interface

This is the most widespread endocrine disorder among women, occurring in up to 15% of women in the general population of reproductive age but in 30% to 75% of women presenting for menstrual irregularities or hirsutism, a referral statistic. The **Rotterdam criteria** run in two steps candidates collapse into one. First, exclude other causes of hyperandrogenism such as congenital hyperplasia, Cushing syndrome and androgen-secreting tumours. Only then diagnose on two of three conditions: oligo-ovulation or anovulation, clinical or biochemical signs of hyperandrogenism, or ultrasound-documented polycystic ovaries, noting that normal-appearing ovaries no longer rule out this diagnosis if the other two criteria are met. The exclusions to order: hyperprolactinaemia and thyroid disorders are associated with menstrual irregularities and should always be ruled out through serum prolactin levels and thyroid function tests, while rapid onset of marked virilisation should raise suspicion of an androgen-secreting tumour.

Mechanistically, gonadotropin-releasing hormone appears insensitive to ovarian steroids, increased luteinising hormone pulse amplitude and frequency drive thecal androgen synthesis, and insulin stimulates thecal androgen production while lowering sex hormone-binding globulin. An elevated luteinising hormone to follicle-stimulating hormone ratio is one of the most common laboratory results, with no numeric threshold given. The metabolic burden is not simply obesity: insulin sensitivity is decreased by 35% to 40% independent of obesity, approximately 50% are obese, and undiagnosed type 2 diabetes is 7 to 10 times more common than in women without it of similar age, with a homeostatic model assessment of insulin resistance greater than 2.3 indicating insulin resistance. The metabolic syndrome criteria belong to the metabolic syndrome, obesity and mortality in serious mental illness notes.

Comorbidity is poorly explained: major depression is commoner than in age-matched controls, a meta-analysis controlling for body mass index finding rates as high as four times greater, yet androgen levels do not consistently differ between depressed and non-depressed women with the syndrome, insulin resistance being a possible common link. Treatment carries a psychiatric tail: metformin decreases insulin resistance, acne, hirsutism and bioavailable testosterone, low-dose oral contraceptives are frequently first-line for acne and hirsutism, and clomiphene citrate is first-line for women who desire pregnancy, while spironolactone may be difficult to tolerate, producing mood changes, decreased libido and fatigue - a profile otherwise read as relapse. One published case used metformin 850 mg per day, titrated to 2,550 mg per day, with spironolactone 100 mg per day: a worked case, not a dose range.

25.9 Valproate and the reproductive endocrine system

in the most extensive study to date of bipolar disorder patients, 10.5% of women on valproic acid developed PCOS compared to 1.4% of women on other mood stabilizers. An independent reference corroborates both figures, gives the phenotype as oligomenorrhoea and

hyperandrogenism, meaning hirsutism, acne, male-pattern alopecia and elevated androgen levels, and supplies the relative risk of 7.5, noting the effect is apparently reversible.

TRAP **The valproate and polycystic ovary syndrome disagreement is about attribution, not observation, and answering it as yes or no is the error.** The same textbook that supplies that figure states elsewhere that the relationship remains controversial because patients with epilepsy and obesity appear to be at increased risk for the syndrome, studies in women with bipolar disorder taking valproate or lithium have reported conflicting results, and in bipolar disorder exacerbation by valproate is less apparent and more inconsistent. The prescribing guidelines take an intermediate line: valproate can cause hyperandrogenism in women and has been linked with the development of polycystic ovaries, the evidence being conflicting. Note the asymmetry: hyperandrogenism asserted as caused, polycystic ovaries only linked. They disagree on causation, not on observation, because obesity, epilepsy itself and pre-treatment menstrual abnormality are independent risks the studies could not separate.

All four converge on conduct. Menstrual history and weight should be followed carefully in women of reproductive potential who begin treatment with valproate; all aspects of the syndrome are reversible in patients with epilepsy, especially if valproate is discontinued and lamotrigine substituted, the effect being apparently reversible in the bipolar series, where no source extends the lamotrigine substitution finding; the syndrome can be mitigated by birth control pills; and this should not preclude women with bipolar illness from treatment with valproate. No source states what to measure or at what interval, so no monitoring schedule exists beyond that instruction. The far larger reason to avoid valproate in women of childbearing potential is teratogenicity, taught with the drug in the Mood disorders: pharmacotherapy notes.

25.10 Prolactin, the pituitary and their psychiatric consequences

Prolactin is produced by the anterior pituitary and inhibited by dopamine from the tuberoinfundibular neurons of the arcuate nucleus. Normal levels run from 5 to 25 ng/mL in women and from 5 to 15 ng/mL in men, and because levels fluctuate during the day, peak during sleep, and rise with exercise and emotional stress, a single raised sample settles nothing. Symptoms usually develop once **prolactin reaches 60 to 100 ng/mL**, with amenorrhoea, gynaecomastia in men, low libido, decreased fertility and galactorrhoea, and headache and visual impairment where there is an adenoma. That is a symptom threshold only: no standard source states a concentration above which a **prolactinoma** becomes likely, so the imaging decision stays qualitative and disproportionate or symptomatic elevation warrants endocrine referral.

The reproductive and pituitary causes, otherwise blamed on a drug: prolactinomas and idiopathic hyperprolactinaemia are commonest, both more common in women. Beyond them:

- Hypothyroidism, since thyrotropin-releasing hormone stimulates prolactin release: the reason thyroid function is checked before the antipsychotic is blamed.
- Acromegaly, through the lactogenic properties of growth hormone, and hypothalamic disease such as craniopharyngioma.
- Pregnancy, breastfeeding and nipple stimulation, and stress, seizures and renal impairment.

The antipsychotic material belongs to the managing antipsychotic metabolic disturbance and hyperprolactinaemia notes. All the SSRIs have case reports, and the one completed study reported that 4.5% of men and 22% of women developed hyperprolactinaemia after 12 weeks of fluoxetine, with oral contraceptives, tricyclic antidepressants, metoclopramide and verapamil also implicated; one study is thin evidence, and that hedge travels with the figure. Sustained elevation suppresses the hypothalamic-pituitary-gonadal axis, producing sexual dysfunction, menstrual disturbance and galactorrhoea, with long-term reductions in bone mineral density. Also, pituitary tumours are found in up to 20% of adults at autopsy, and as many as 30% are non-secretory.

PITFALL

Simple, non-formed visual hallucinations in a patient who retains insight are treated as psychosis and answered with an antipsychotic. Visual hallucinations have been reported where a pituitary tumour impinges the optic chiasm, are more often simple and non-formed, are associated with insight, and can be exacerbated by bromocriptine; the accompanying sign is a field defect, since large pituitary tumours produce the classic bitemporal hemianopia, oculomotor palsies and headache. Confrontation fields take a minute and are almost never done in a psychiatric clerking.

Bromocriptine normalises serum prolactin in 70% to 80% and decreases tumour size in more than 50% of patients with prolactinomas. Cabergoline is the preferred drug for hyperprolactinaemia, initiated at 0.25 mg twice a week or 0.5 mg once a week; the pharmacology text's escalation sentence is truncated mid-clause, so no maximum dose is quoted. Medication must continue indefinitely, since discontinuation usually leads to relapse. The hazard runs both ways, since psychosis has been described on bromocriptine for a prolactinoma, and impulse control disorders such as pathological gambling with cabergoline.

25.11 Androgens, anabolic steroids and mood

In general, 200 mg of testosterone cypionate taken every 2 weeks will restore physiologic testosterone concentrations in a hypogonadal male, and illicit use escalates beyond that because a eugonadal male given physiologic dosages gains nothing, exogenous steroid shutting down endogenous production via feedback inhibition of the hypothalamic-pituitary-testicular axis. So illicit users take dosages 10 to 100 times greater than therapeutic doses, while transdermal gels or patches are generally inadequate to produce massively supraphysiologic levels. Testosterone and hypogonadism sit at the opposite end of the same axis: where it is genuinely deficient, testosterone improves mood and decreases irritability in hypogonadal men.

WEEKLY TESTOSTERONE-EQUIVALENT EXPOSURE	PSYCHIATRIC EFFECT REPORTED	WHAT THE BAND IS GOOD FOR
300 mg of testosterone a week or less	Psychiatric effects appear rare	Why negative laboratory trials do not reassure about illicit use.
Between 300 and 1,000 mg of testosterone equivalent per week	Mood syndromes appear more frequent, albeit still uncommon	Where most gym-population field studies sit.
More than 1,000 mg a week	Mood syndromes become quite common and are occasionally severe; described psychotic cases are primarily in this band	The exposure to ask about in a young muscular man presenting manic or paranoid.

These are exposure bands, never a prescribing range. Across four blinded placebo-controlled studies giving at least 500 mg a week, 5 of 109 men exhibited prominent hypomanic or manic syndromes, and that crude rate of approximately 5% may represent a lower bound for illicit users, who stack drugs. Psychotic symptoms usually consist of grandiose or paranoid delusions within a manic episode, disappearing within a few weeks of discontinuing the agent.

Withdrawal is where a psychiatrist usually meets this. Because these agents suppress hypothalamic-pituitary-testicular function, male users often become hypogonadal after stopping a course; function usually recovers within a few weeks to a few months, but hypogonadism may persist for more than a year and appears partially or completely irreversible in some cases. Withdrawal depression is usually brief and self-limited, but some episodes have continued for months and required fluoxetine, tricyclic antidepressants or ECT, and completed suicides on withdrawal have been reported. The loop is self-sustaining: stopping produces hypogonadism and depression, the user resumes to self-treat it, and some develop brittle mood disorders with mixed symptoms and prominent impulsivity - the authors' clinical experience, and the differential for a young man with a rapid-cycling mixed state.

PITFALL

A testosterone result is read against a memorised cut-off rather than the reporting laboratory's reference interval, and one textbook shows why that fails. One chapter gives normal testosterone plasma concentrations for men as 300 to 1,000 ng/dL. Another, reporting former users, classifies against a normal reference limit of 12.1 nanomoles per liter, 350 nanograms per deciliter. The conversion is internally correct, so the disagreement about where normal starts is genuine, most plausibly because the second is one laboratory's limit. Neither is a diagnostic threshold for hypogonadism, and these sources state none.

GOOD TO REMEMBER

Recovery is sequenced and endocrinologist-led, though these sources give no doses for it. Physiologic doses of exogenous testosterone first, after supraphysiologic doses stop; then clomiphene, stimulating the pituitary to resume luteinising hormone and follicle-stimulating hormone production; simultaneously human chorionic gonadotropin to stimulate the testis; then a gradual taper as the axis reawakens. The point is relapse prevention, since withdrawal symptoms untreated will strongly tempt a man to resume use.

No assay makes any diagnosis in this topic: premenstrual dysphoric disorder is confirmed by prospective daily ratings during at least two symptomatic cycles, the perimenopause is staged on menstrual pattern, and a raised prolactin or low testosterone is read against the reporting laboratory's interval, never a remembered number.

26 · Toxic, occupational and environmental encephalopathies

These are the encephalopathies whose cause is written on a payslip, a water source or a bottle, and never on the referral letter. A battery-factory worker, irritable, forgetful and constipated, is referred with depression; a farmer arrives sweating and confused and nobody mentions the tin in the shed. The diagnosis turns on an exposure question the history never reaches.

Four axes make the agents memorable: route of entry, anatomical target, assay and reversibility. Heavy-metal toxicity is where those axes separate most cleanly, because lead, mercury, arsenic and manganese differ on every one of them at once. Delirium as a syndrome belongs to the Stroke, TBI, delirium and focal brain syndromes notes; the dementia syndrome and its workup to the Dementias and neurocognitive disorders notes; antipsychotic movement disorders to the Schizophrenia: management, antipsychotics and adverse effects notes.

26.1 Lead: how it gets in, where it goes, and what it does to a neurone

Gut absorption varies with age and diet, children absorbing about **40%** of ingested lead against less than 20% in adults, and fasting, calcium deficiency or iron deficiency increases it further. Inhaled lead is absorbed at about 90%, and tetraethyl lead crosses skin while inorganic lead does not.

A blood level maps the body burden poorly. Lead binds haemoglobin, distributes to renal tubular epithelium and liver, then redistributes to bone, teeth and hair, so that **95%** of the adult body burden sits in bone. Growing bone accumulates more and can form **lead lines** on radiography. Bone lead is slowly reabsorbed, faster when calcium is depleted, including in pregnancy. Only small quantities reach brain, mostly grey matter and basal ganglia, and lead crosses the placenta.

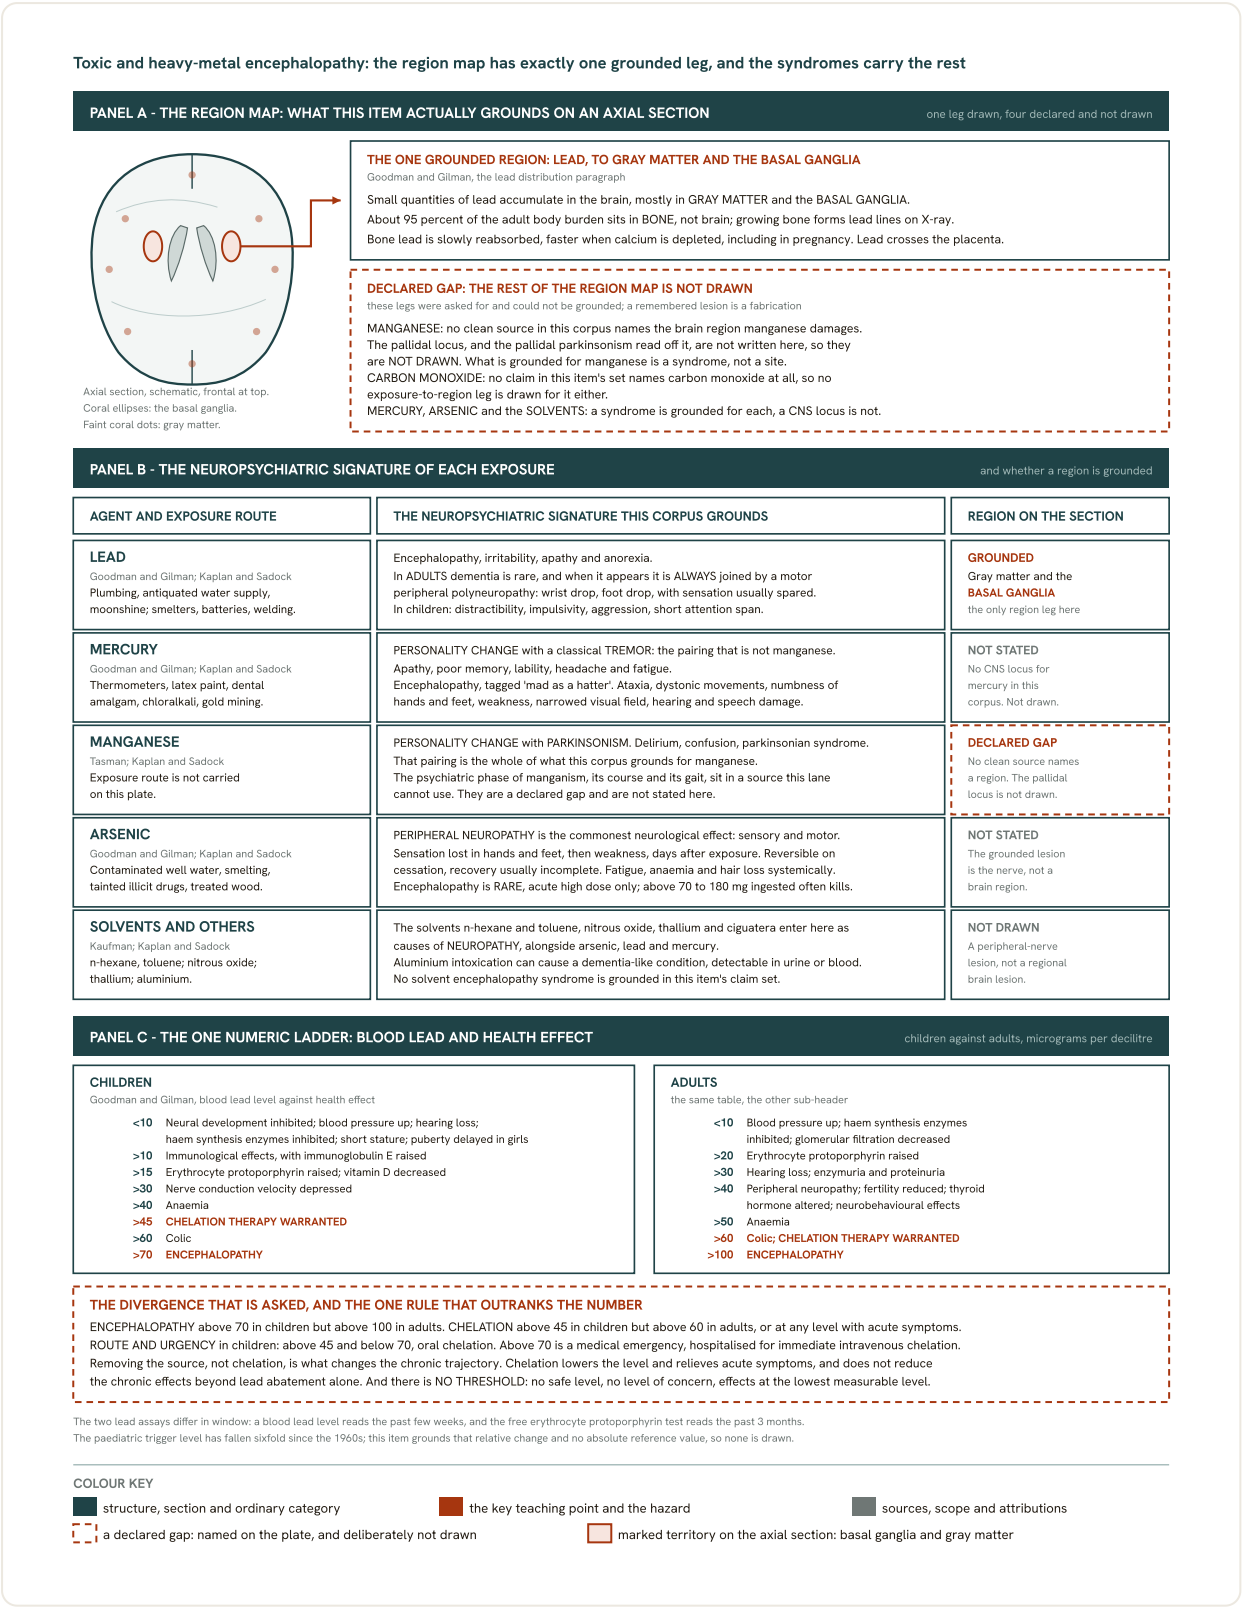


Fig 33.10 The toxic and heavy-metal region map, drawn only as far as the sources go: one grounded anatomical leg, the declared gaps named beside it, and the syndromes that carry the teaching instead. Lead is the only exposure this item maps to tissue, to gray matter and the basal ganglia; the manganese-to-pallidum and carbon monoxide legs are named on the plate and deliberately not drawn. The differentiating triad is syndromic: manganese gives personality change with parkinsonism, mercury personality change with a tremor, and adult lead intoxication only rarely a dementia, then always with a motor polyneuropathy. (Goodman and Gilman's The Pharmacological Basis of Therapeutics; Tasman Psychiatry; Kaplan and Sadock's Comprehensive Textbook of Psychiatry 2024; Kaufman's Clinical Neurology for Psychiatrists.)

Lead injures by **molecular mimicry** of zinc and calcium, altering protein structure. Centrally, inhibition of calcium channels and of calcium-responsive proteins such as protein kinase C and calmodulin drives inappropriate transmitter release across dopaminergic, cholinergic and glutamatergic pathways, while high concentrations disrupt the blood-brain barrier. Encephalopathy is the same lesion with a failing barrier added.

26.2 The blood lead ladder, and why no level is reassuring

>45	>70	>60	>100
UG/DL, CHELATION WARRANTED, CHILDREN	UG/DL, ENCEPHALOPATHY, CHILDREN	UG/DL, CHELATION WARRANTED, ADULTS	UG/DL, ENCEPHALOPATHY, ADULTS

Blood lead concentration and health effect, children against adults

EFFECT	CHILDREN (UG/DL)	ADULTS (UG/DL)
Present at the lowest measurable levels	Below 10: inhibited neural development, raised blood pressure, hearing loss, inhibited haemoglobin synthesis enzymes, short stature, delayed sexual maturation in girls	Below 10: raised blood pressure, inhibited haemoglobin synthesis enzymes, reduced glomerular filtration
Immunological effects (raised IgE)	Above 10	Not listed
Raised erythrocyte protoporphyrin	Above 15	Above 20
Reduced vitamin D	Above 15	Not listed
Hearing loss	Below 10	Above 30
Enzymuria and proteinuria	Not listed	Above 30
Depressed nerve conduction velocity	Above 30	Not listed
Peripheral neuropathy	Not listed	Above 40
Reduced fertility, altered thyroid hormones, neurobehavioural effects	Not listed	Above 40
Anaemia	Above 40	Above 50
Chelation therapy warranted	Above 45	Above 60
Colic	Above 60	Above 60
Encephalopathy	Above 70	Above 100

At the bottom of the ladder there is no line at all: lead disturbs synaptic pruning, neuronal migration and neurone-glia interaction, producing falls in IQ, language and motor deficits, distractibility, impulsivity and aggression, with no evidence of a threshold and associations evident at the lowest measurable levels.

RULE **A blood lead level can confirm poisoning; it can never exclude harm.** Regulation has followed the biology: there is no safe level of lead exposure for children, the CDC no longer provides a level of concern, and the emphasis has moved to primary prevention. The paediatric trigger level has fallen sixfold since the 1960s, an administrative concession rather than a threshold.

26.3 The adult lead patient, and the management of metal poisoning

Adults present differently. **Lead palsy**, wrist drop and foot drop from degeneration of motor neurones, was classical in painters and is now rare. Tasman gives the discriminator: lead intoxication in adults only rarely produces a dementia, and when it does it is always accompanied by a motor peripheral polyneuropathy with wrist or foot drop. Exposure is inhalational and trade-specific: greatest in smelter and storage-battery workers, with further risk in welding, construction, printing and radiator repair, to which Kaplan and Sadock add plumbing trades, an antiquated water supply and illegal moonshine.

Management of any heavy-metal toxicity turns on one uncomfortable fact: the chelator treats the number, removal from exposure treats the disease. LOCUS is set by concentration and symptoms: chelation is warranted above 45 in children and above 60 ug/dL in adults, or where there are acute symptoms of lead poisoning, and oral chelation is recommended for a child above 45 but below 70 ug/dL, and above 70 is a medical emergency requiring hospitalisation and immediate intravenous chelation, on the American Academy of Pediatrics position. That is the only guideline-level statement available here. FOCUS by phase: lower the burden, then keep the patient away from the workplace or water source. MODUS is narrow. Removing the source is the most important response; chelation lowers the blood level and relieves acute symptoms but does not reduce the chronic effects of lead beyond abatement alone, and in rats chelators mobilise lead into the brain.

Chelator selection is a matching exercise: edetate for acute lead poisoning; dimercaprol for acute arsenic, gold and mercury poisoning and for acute lead poisoning combined with edetate; succimer for children with elevated blood lead; penicillamine for copper; trientine for Wilson disease. Copper chelation is taught where this chapter deals with Wilson's disease. Succimer is preferred because it is orally bioavailable and, being hydrophilic, does not mobilise metals to the brain, and it does not significantly chelate essential metals such as zinc. Dimercaprol is given by deep intramuscular injection as a 100 mg/mL solution in peanut oil, contraindicated in peanut allergy. A formulation strength, not a dose: no milligram-per-kilogram regimen exists in these sources.

PITFALL

Reaching for the intravenous route in the sickest lead-poisoned patient because they are the sickest. Edetate is principally nephrotoxic, so adequate urine output must be established before and during treatment; and in lead encephalopathy with cerebral oedema the intravenous route can produce a potentially fatal rise in intracranial pressure, so the intramuscular route is preferred.

26.4 Arsenic: well water, neuropathy and arsine gas

Arsenic's signature is peripheral, not central, and the route that matters in South Asia is drinking water: massive ingestion, smelting of lead and copper ore, contaminated well water, tainted illicit drugs and the burning of preserved wood. Goodman and Gilman add: the commonest effect of acute or subacute exposure is a peripheral neuropathy involving both sensory and motor neurones, appearing several days after exposure and reversible on cessation though recovery is usually incomplete; encephalopathy is rare and follows only acute high-dose exposure. Arsenic is a poor answer to a chronic confusional state.

TRAP **Merging the two standard descriptions of arsenic neuropathy into one sentence.** Goodman and Gilman describe loss of sensation in the hands and feet, often followed by muscle weakness, a negative sensory phenomenon in both territories at once. Kaplan and Sadock's table describes burning and numbness of the feet followed by the hands, a positive, painful, length-dependent ascent from the feet. Hold both: the narrative account is the stronger because the table cell is positionally reconstructed. Answer the picture the stem describes.

Chronic vascular disease is the other signature. **Blackfoot disease** is cyanosis of the extremities progressing to gangrene, endemic in regions of Taiwan where well water contains arsenic at 170 to 800 ug/L, a water concentration and not a human threshold. Ingestion above 70 to 180 mg is often fatal, immediate death coming from cardiac and gastrointestinal effects. **Arsine gas**, formed industrially by reduction of arsenic, causes rapid and often fatal haemolysis, with haemoglobinuria and anuria hours later and jaundice after 24 hours; roughly **25%** of exposures are fatal. Treatment splits acute from chronic: chelation works after short-term exposure but has little or no benefit in the chronically exposed, and exchange transfusion is often warranted after arsine exposure.

26.5 Mercury and manganese, and separating the metals at the bedside

Mercury is a three-syndrome metal because it exists in three forms: elemental mercury, met as vapour; inorganic mercurous and mercuric salts; and organic compounds, of which **methyl mercury**, generated environmentally by aquatic microorganisms, is the one of concern. Covalent bonding with sulfhydryl residues at very low concentrations accounts for most biological effects. The form predicts the target. Inhaled elemental vapour is acutely toxic to lung, progressing to interstitial pneumonitis. The salts strike the kidney: acute exposure, typically in suicide attempts, produces renal tubular necrosis, while chronic exposure causes predominantly glomerular injury. Occupationally, the chloralkali industry carries the highest risk, with exposure also in dentistry and gold mining. The index case is Minamata, where fish contaminated with methylmercury caused severe learning disability in Japan in the 1950s.

Separating the metals by neuropsychiatric signature and confirmatory specimen

METAL	NEUROPSYCHIATRIC SIGNATURE	DETECTED IN
Lead	Dementia only rarely, always with a motor polyneuropathy; encephalopathy, irritability, apathy, anorexia	Blood lead level for the past few weeks; free erythrocyte protoporphyrin for the past 3 months
Mercury	Personality change with a classical tremor; apathy, poor memory, lability, headache; encephalopathy tagged mad as a hatter, with ataxia, dystonic movements and narrowed visual fields	Urine, blood, hair
Manganese	Personality change with parkinsonism; delirium, confusion, parkinsonian syndrome	Urine, blood, hair
Arsenic	Fatigue, loss of consciousness, anaemia, hair loss	Urine, blood, hair
Aluminium	A dementia-like condition	Urine or blood

Personality change is shared by manganese and mercury, separated only by the motor sign beside it, parkinsonism against tremor. The assays are asymmetrical: lead has a quantified ladder and a second test with a longer memory, the others only a specimen. So is the anatomy. Lead has a stated locus, but **manganism** is recorded only as parkinsonism produced by high-level manganese exposure, with no locus stated anywhere here, so the accompanying figure carries no grounded manganese leg. Kaufman places heavy-metal toxicity among the toxic causes of neuropathy, beside ciguatera fish poisoning, thallium, nitrous oxide, n-hexane and toluene.

GOOD TO REMEMBER

A worker removed from a smelter a month before presentation can show a falling blood lead level beside a still-abnormal protoporphyrin. The pair dates the exposure, which matters when the question is whether the workplace poisoned them, not whether they are poisoned now.

26.6 Carbon monoxide: the gas that leaves a hypoxic footprint

Carbon monoxide binds haem tightly to form **carboxyhaemoglobin**, which can no longer carry oxygen. Smokers of one to two packs daily typically run 3% to 8% and non-smokers under 2%; after acute poisoning 5% to 10% is associated with impaired alertness and cognition, and **30% to 50%** causes significant falls in tissue oxygen delivery. The smoker's baseline overlaps the symptomatic band, so a single carboxyhaemoglobin figure cannot grade the severity of carbon monoxide poisoning in a smoker.

Poisoning usually arises from deliberate self-harm with car exhaust and, more recently, from charcoal burning in the Far East, though it also follows badly ventilated boilers and fires. anoxia, hypoglycaemia and carbon monoxide produce similar patterns of injury, in which cerebral and cerebellar atrophy may occur and the hippocampus and globus pallidus are particularly

vulnerable. The lesion is hypoxic rather than gas-specific: a hypoglycaemic coma and a charcoal-burning attempt leave the same deficits, and the amnesia follows hippocampal vulnerability.

The recovery trajectory is relatively brief compared with traumatic brain injury; parkinsonism is not infrequent after a latent interval; the commonest cognitive impairment is poor memory. The latent interval is the examinable phenomenon: the patient improves, is discharged, then deteriorates. No numeric latency is given. Nor is a complication rate: the prevalence of complications is unclear, with a wide range of figures reported. That refusal is the citable fact.

26.7 Organophosphates: the toxidrome, ageing, and the emergency

Organophosphate exposure is the toxic encephalopathy an Indian psychiatrist meets most often, and the only one here in which minutes matter. Organophosphate insecticides bind and inactivate acetylcholinesterase, so acetylcholine accumulates and irreversibly depolarises postsynaptic neuromuscular junctions; people committing suicide, especially in India, often deliberately drink organophosphate pesticides, and the nerve gases including sarin act identically.

Goodman and Gilman organise the syndrome into three compartments. Acute intoxication produces muscarinic and nicotinic features and, except with quaternary compounds of low lipid solubility, central signs; effects appear within minutes of inhaled vapour but are delayed after gastrointestinal and percutaneous absorption. Route decides which system declares first: vapour brings ocular and respiratory features first, with marked miosis, ciliary spasm and brow ache; ingestion brings gastrointestinal features earliest; percutaneous exposure produces localised sweating and fasciculation at the site.

MNEMONIC

DUMBBELLSS

Diarrhoea, urination, miosis, bronchorrhoea, bradycardia, emesis, lacrimation, lethargy, salivation, sweating: the muscarinic effects of the **cholinergic toxidrome**. Against the anticholinergic row of the same table the discriminators are three paired opposites: mydriasis against miosis, decreased against increased bowel sounds, decreased against increased diaphoresis.

The nicotinic and central compartments carry the mortality. Nicotinic action at the neuromuscular junction gives weakness, fasciculation and finally paralysis, gravest at the respiratory muscles; central effects include confusion, ataxia, convulsions, coma and central respiratory paralysis. Death ranges from under 5 minutes to nearly 24 hours and is primarily respiratory. Respiratory failure arrives by three converging routes, so airway management outranks every antidote.

LOCUS is the emergency department and then intensive care, never the psychiatric ward; assessment of suicidal intent waits until the crisis is over. FOCUS by phase is decontamination and airway, then muscarinic blockade and enzyme reactivation, then the later syndromes. MODUS begins before any drug: terminate exposure, remove contaminated clothing, wash skin copiously or lavage the stomach, maintain a patent airway, give oxygen, and alleviate persistent

convulsions with diazepam 5 to 10 mg intravenously. For muscarinic effects, atropine is given in doses sufficient to cross the blood-brain barrier, beginning with 2 to 4 mg intravenously, then **2 mg every 5 to 10 minutes** until muscarinic symptoms disappear or signs of atropinisation appear.

Atropine alone is not enough. It antagonises the muscarinic actions but is virtually without effect against the peripheral neuromuscular compromise, which pralidoxime reverses, the oxime acting most markedly at the skeletal neuromuscular junction. No pralidoxime dose is printed; none is recoverable in clean form here. The urgency rests on mechanism: stability of the phosphorylated enzyme is enhanced by **ageing**, the loss of one alkyl group, after which recovery depends largely on synthesis of new enzyme protein. An aged enzyme cannot be reactivated by any oxime, making the oxime a race rather than a routine.

PITFALL

Excluding organophosphate poisoning because the pupils are not small. With acute systemic absorption, miosis may not be evident because of sympathetic discharge in response to hypotension. The sickest patient is the one in whom the cardinal sign disappears.

The later syndromes are named but not defined. Kaplan and Sadock list among organophosphate central effects emotional lability, reduced alertness, cognitive impairment, convulsions and coma, with weakness of neck flexors and proximal muscles, and **organophosphate-induced delayed polyneuropathy**. Goodman and Gilman record only that an intermediate syndrome is recognised and that delayed neurotoxicity and recurrent seizures may follow severe intoxication. Latencies and the enzyme behind the delayed polyneuropathy are stated nowhere in these sources.

IN INDIA

These three exposures are the common case here, not the exotic one. The cholinergic toxidrome is a working diagnosis in any rural presentation combining altered sensorium with sweating, secretions and fasciculation. A neuropathy in a patient from a tube-well belt is a toxicological question, not only a diabetic one. A cluster of drinkers with gastroenteritis, delirium and blurred vision is a methanol outbreak until proved otherwise. Laboratory support is largely absent, so treat on the toxidrome and the atropine response, and ask about the water source, the workplace and the liquor as routinely as about mood.

26.8 Solvents and inhalants: toluene, the hexacarbons and carbon disulfide

Solvent encephalopathy divides by route into occupational and deliberate exposure, and the deliberate does far more damage. The hexacarbons are constituents of glues, lacquers and solvents, also inhaled intentionally, producing central depression or stupor acutely and, on repeated glue sniffing, cognitive deficits, motor and sensory neuropathy and autonomic dysfunction; treatment is cessation of exposure. That last clause is the prognosis: no antidote exists.

The systemic toll of chronic inhalant use spans rhabdomyolysis, renal disturbance, cardiomyopathy, hepatotoxicity, pulmonary disease and haematopoietic disorders including methaemoglobinaemia, aplastic anaemia and acute myelocytic leukaemia, with peripheral neuropathy, reversible cerebellar signs or frank cerebellar degeneration, widened sulci and basal cisterns on imaging, and two distinct dementias: lead encephalopathy from leaded petrol and **white matter dementia** from toluene. Tasman's catalogue is headed, importantly slowly reversible and/or irreversible.

Toluene is the solvent most often examined. Most reports emphasise combined cerebral and cerebellar impairment with pyramidal change, ranging from mild cognitive impairment to severe dementia with cerebellar ataxia, oculomotor abnormality, tremor, deafness and hyposmia; cognitive dysfunction is the most frequent and most disabling feature and may be the earliest sign of permanent damage. It is not the reversible fog of intoxication clinicians discount. Magnetic resonance imaging in chronic abusers shows diffuse central white matter change comprising

- cerebral, cerebellar and brain stem atrophy;
- loss of grey-white differentiation throughout the central nervous system;
- increased periventricular white matter signal on T2-weighted images.

Set that against carbon monoxide, which strikes grey matter nuclei, and methanol below. Carbon disulfide completes the set: an industrial polar solvent affecting factory workers and a major metabolite in the breakdown of disulfiram.

26.9 Toxin and non-antipsychotic medication-induced movement disorders, and methanol

Several agents causing parkinsonism are already taught above, so the toxin-induced movement disorders are largely a second pass over exposures this chapter has already met. DSM-5-TR names as toxic causes MPTP, organophosphate pesticides, manganese, methanol, cyanide, carbon monoxide and carbon disulfide. Brazis adds disulfiram, toluene, n-hexane and welding-related parkinsonism on the toxin side and reserpine, tetrabenazine, lithium, amiodarone, diltiazem and ethanol on the drug side. Kaplan and Sadock note that associations with well water, farming and rural living may reflect neurotoxins in pesticides containing homologues of MPTP. Carbon monoxide and manganese also cause secondary dystonia, beside levodopa, anticonvulsants, flecainide and ergotamines.

TRAP **Reading welding-related parkinsonism in a list as settled aetiology.** Three standard sources list manganese, one of them welding, flatly among toxin causes. Kaufman disputes the epidemiology: the relationship between welding or other manganese exposure and the risk of developing Parkinson disease is controversial, and although these were previously considered significant risk factors, more recent epidemiological studies have not indicated increased risk. His next sentence concedes that high levels of manganese, that is manganism, lead to parkinsonism. All four agree that high-level exposure produces a parkinsonian syndrome; what they dispute is whether ordinary occupational welding raises the risk of Parkinson disease itself.

Whether manganese parkinsonism improves after removal from exposure, or responds to levodopa, is stated nowhere in these sources: the one entity here whose reversibility cannot be answered. MPTP, a by-product of the synthesis of a certain opioid, kills dopaminergic cells and induces permanent parkinsonism, and it selectively poisons nigrostriatal tract neurones.

Two non-antipsychotic problems are missed for one reason, latency. DSM-5-TR lists as non-antipsychotic causes of medication-induced parkinsonism the calcium channel antagonists flunarizine and cinnarizine, the dopamine depleters reserpine and tetrabenazine, the antiepileptics phenytoin, valproate and levetiracetam, antidepressants, lithium, chemotherapeutic agents and the immunosuppressants ciclosporin and tacrolimus. Onset is unreliable: parkinsonism usually develops a few weeks after starting or raising the dose, but may develop insidiously after many months, and mainly with calcium channel blockers a second peak of onset is reported after about 1 year. Tasman supplies the longest latency: valproic acid may, albeit very rarely, cause a combination of parkinsonism and dementia, and because this evolution is delayed for from 6 months to 4 years the connection may be missed. It is also the closest these sources come to a medication-induced encephalopathy, since the other non-antipsychotic drugs named here are credited with a movement disorder alone. Lithium as a drug belongs to the Mood disorders: pharmacotherapy notes.

> In a new parkinsonian or cognitive syndrome, take a drug and exposure history covering the preceding several years, not the preceding several weeks.

Methanol has the tightest anatomy here and the most preventable outcome. Alcohol-dependent patients may inadvertently drink methanol, a component of antifreeze and cooking fuels and an illicit adulterant of everyday ethanol, producing gastroenteritis, delirium and visual disturbance, particularly blurred vision and a scotoma; with severe intoxication or merely chronic intermittent consumption the optic nerves atrophy and victims go blind. Methanol preferentially damages the optic nerves and also the putamen, where it may cause parkinsonism. Management has three limbs: bicarbonate for the acidosis; ethanol to divert the metabolism that yields formaldehyde and formic acid; and haemodialysis, with fomepizole possibly more effective than ethanol. No dose for any of the three is stated in these sources. Kaplan and Sadock list arsenic, lead, N-hexane, organophosphates, pyridoxine, cisplatin, vincristine and dapsone among environmental and medication toxicities, where pyridoxine is the reminder that a vitamin in excess belongs in a toxicology chapter.

Consolidate and apply

27 · Worked vignettes

A vignette in this territory is written so that the psychiatric syndrome is easy and the body is not. The stem hands over a recognisable picture of depression, psychosis, anxiety, confusion or personality change, then buries the discriminating evidence in items that read like background: a bowel habit, a skin change, a drug bought without prescription, an occupation, a value quoted in

passing. Candidates who read for the psychiatric label find it within two lines and elaborate a diagnosis the examiner never doubted. The marks sit in the items that were skipped.

These cases consolidate the chapter rather than replacing its disease sections. Each is followed by the reasoning that should be audible in a written answer or a viva: the syndrome, the system that could have produced it, the finding that most changes the ranking, and what must happen before the patient leaves. Criteria, reference ranges, thresholds, doses, correction rates and monitoring intervals belong to the sections that own them.

27.1 Read the stem twice: once for the syndrome, once for the body

Read these stems in two deliberate passes. The first names the psychiatric syndrome exactly as presented, without interpretation. The second ignores the mental state entirely and collects everything somatic: physical symptoms and examination findings, every drug prescribed and unprescribed, exposures, diet, reproductive and family history, tempo, and any value already quoted. Only then ask which collected item would alter the diagnosis if it were deleted from the stem.

Syndrome	System	Direction	Action
WHAT IS PRESENTED	WHAT COULD PRODUCE IT	WHICH WAY CAUSE RUNS	WHAT CANNOT WAIT

The **diagnostic pivot** is the finding, or cluster of findings, that most alters the ranking of explanations, and in this chapter it is almost never a psychiatric feature. It is a physical sign, an exposure, a temporal coincidence with a prescription, or a discordance between the psychiatric diagnosis and the age, tempo or context of onset. Nominate it explicitly: a stated pivot converts a list into an argument.

The **opening formulation** should then sit at the level the evidence supports. Lead with the syndrome, add the system you believe generated it and the strength of that belief, name the strongest competitor, and state the observation that would separate them. Naming a specific disease in the first sentence is a gamble that pays only when the somatic signature is unambiguous, and it costs the answer when the stem was built around a different organ.

RULE **Every answer here names three things: the psychiatric syndrome, the bodily system that could have produced it, and the investigation that separates them.** Missing the second reads as a psychiatric formulation that happened to be asked in a medical setting. Missing the third asserts a cause it has not tested, which is the failure this chapter exists to prevent.

27.2 Vignette: the depression that will not lift, in a body that has changed

A middle-aged woman is referred with several months of low mood, loss of interest, profound fatigue and poor concentration. Successive adequate antidepressant trials have not helped. Her family describe her as slowed rather than sad, and mention weight gain, constipation, dry skin, hair thinning and cold intolerance, all attributed to the depression. Nobody has examined her since the first referral.

This is the chapter's commonest stem shape: a **psychiatric phenocopy**, a systemic disease presenting through a psychiatric door with a mental state indistinguishable from primary illness.

The pivot is not the mood but the convergence of physical changes belonging to one axis. Each item alone is explicable as a depressive or medication effect; together, and in one direction, they point at an organ.

Apparent treatment resistance is the second pivot. Failure to respond to adequate treatment is a standing trigger for looking beyond the psychiatric diagnosis, and the correct response is to re-examine the patient rather than re-read the drug chart. Say what would be sought on examination, order the test the somatic signature calls for, and state that reference ranges and the interpretation of a borderline result belong to the thyroid sections.

The management answer has two halves and candidates write only the first. Correct the systemic disease with the physician who will manage it, expecting the psychiatric syndrome to improve alongside. Then reassess: a residual affective picture persisting after biochemical correction is a legitimate target for independent psychiatric treatment, not an argument that the endocrine finding was irrelevant. Hormone augmentation and the endocrine consequences of mood-stabilising treatment are taught in their own sections; the drug itself, with its pharmacokinetics and non-thyroid organ effects, belongs to the Mood disorders: pharmacotherapy notes.

GOOD TO REMEMBER

When a depression has failed adequate treatment, the next investigation is a physical examination, not another antidepressant. Weigh the patient, look at the skin, hair and neck, take the pulse, check for a postural drop and elicit the reflexes. A stem saying nobody has examined the patient is telling you where the answer is.

27.3 Vignette: the new confusion in a patient who already has a diagnosis

A patient with established severe mental illness is admitted for a systemic medical problem and referred back to psychiatry for agitation and confusion. The picture fluctuates: lucid in the morning, disorganised and drowsy by evening, with disturbed sleep, variable attention and an abnormal movement of the outstretched wrists that the encephalopathy sections teach as organ-localising. The referral asks for sedation and says the patient has relapsed.

The correct opening is an **acute metabolic encephalopathy** until proved otherwise. Three features carry it: the change is new against the patient's own baseline, conscious level and attention fluctuate rather than thought content alone, and there is a physical sign no psychiatric illness produces. The disease that generated the state and its correction are this chapter's material. The acute confusional syndrome itself, with its criteria, structured assessment, workup and prevention, belongs to the Stroke, TBI, delirium and focal brain syndromes notes.

Make the search for the driver audible: which organ of clearance is failing, which solute has moved, which drug was started or stopped around the change, and whether the illness, its treatment, hypoxia, sepsis, deficiency or withdrawal is responsible, singly or together. Sedation before that search is the harm: it removes the fluctuating conscious level that was the diagnostic evidence, and can deepen the encephalopathy already present.

TRAP **Letting an existing psychiatric diagnosis absorb a new deterioration.** Candidates read the psychiatric history first and file every later change under it, so confusion becomes relapse, drowsiness becomes negative symptoms and a new physical sign becomes a medication effect. The discipline is the opposite: a patient with chronic mental illness who changes acutely has an acute problem until that is excluded, and the existing diagnosis makes such a presentation more likely, not less.

27.4 Three shorter stems: when to stop and think metabolic

A young adult presents with a first psychotic episode accompanied by a movement abnormality, deteriorating handwriting, mildly deranged liver enzymes and a sibling who died undiagnosed. A second has recurrent episodes combining severe abdominal pain, autonomic instability, weakness and florid psychiatric disturbance, each following a newly started drug and resolving between attacks. A third, an older adult on a restricted diet, has cognitive decline with sensory disturbance in the feet and an unsteady gait.

These test the **inherited metabolic differential** and its nutritional neighbours, and share one feature: the psychiatric syndrome is ordinary and the accompanying findings are not. Name the pattern, then hand the confirmatory test to the section that owns it: copper studies and the ocular sign to the copper-metabolism sections, the correctly timed and handled specimen to the porphyria section, the deficiency assays to the nutritional sections. The marks lie in the decision rule, not the disease names. Stop and think metabolic when the age of onset is wrong for the psychiatric diagnosis, when psychiatric, neurological and systemic features coexist without a unifying explanation, when episodes are stereotyped and drug-triggered, when there is consanguinity or undiagnosed neuropsychiatric illness in the family, or when the course does not behave like its label. Cognitive impairment from these diseases is described here in clinical language and its reversibility is this chapter's business; the dementia syndrome and its workup belong to the Dementias and neurocognitive disorders notes, and cognitive instruments to the Psychotherapies, clinical assessment, MSE and nosology notes.

STEM PATTERN	OPENING FORMULATION	THE PIVOT	WRONG TURN
Depression with convergent somatic change	Depressive presentation of systemic disorder	Findings pointing to one axis	A third antidepressant, not an examination
Anxiety with paroxysmal autonomic surges	Anxiety needing a secretory cause excluded	Stereotyped disproportionate episodes	Treating panic without asking why it began
Fluctuating confusion in known mental illness	Acute metabolic encephalopathy, cause unknown	Fluctuating conscious level plus a sign	Calling it relapse and sedating it
First psychosis with neurological and hepatic signs	Young-onset psychosis, possibly metabolic	Three systems, one young patient	Starting antipsychotics, stopping the workup
Stereotyped attacks after a new drug	Recurrent neurovisceral episode, drug-triggered	Episodicity with a drug relation	Diagnosing somatisation from normal scans
Cognitive decline with a nutritional history	Potentially reversible cognitive syndrome	Diet and neurological signs together	Assuming irreversibility from age
Derangement in a treated psychiatric patient	Treatment-attributable metabolic disturbance	Temporal relation to the prescription	Stopping the effective drug reflexively

27.5 Vignette: the patient whose treatment produced the disease

A young man treated effectively for a psychotic illness for several years has gained substantial weight, with an increase in waist circumference his family notice more than he does. Recent bloods show abnormal glucose and lipid results. His father has diabetes. He is mentally well, wants to stop the medication because of the weight, and has already halved his own dose.

The direction of causation is inverted here and the psychiatrist is the exposure rather than the detective. The reasoning turns on the **attribution question**: how much belongs to the drug, how much to the illness with its associated inactivity, diet, tobacco and social disadvantage, and how much to the baseline that preceded treatment and was never documented. Attribution decides whether changing the drug would change the outcome, and an honest answer concedes that these contributions cannot be cleanly separated in an individual.

Ownership discipline separates a good answer from a long one. The receptor-mediated mechanism by which antipsychotic treatment produces weight and metabolic change, with the comparative liability of agents, is taught in the treatment-burden section. The criteria and thresholds defining metabolic syndrome are stated once, in the metabolic syndrome and mortality section, and the monitoring parameters and their cadence once, in the monitoring section. Choosing between agents belongs to the Schizophrenia: management, antipsychotics and adverse effects notes.

The decision on top of all four has a shape. Take the covert dose reduction as the most urgent item in the stem, because relapse risk is immediate and metabolic risk is not. Treat the complaint as legitimate rather than as poor insight, since the harm is real and was prescribed. Establish the

baseline and what has changed, address modifiable contributors actively rather than advising them, then consider with the patient whether the regimen can be altered, whether a metabolic agent is indicated, and what will be monitored to know whether any of it worked.

27.6 Vignette: the drowsy patient with a low sodium

A patient on established psychotropic treatment becomes nauseated, unsteady and progressively drowsy over a short period, with headache and a witnessed seizure. The serum sodium is low. A second stem describes a long-stay patient who drinks constantly, whose weight rises measurably through the day, who is intermittently confused in the evenings, and whose sodium is also low.

The **discordant fact** in each is the route by which the sodium fell, and it is the whole question. One patient has retained water because a drug altered water handling; the other has taken in more than the kidney can excrete. The history separates them long before any calculation: what was prescribed and when, against what is being drunk and how far the weight moves within one day. Both then converge on the same urgency, because the neurological state and the speed of the fall determine how ill the patient is.

Two handovers keep this answer honest. Mechanisms of drug-induced water retention, the psychotropic classes responsible, the risk groups and the monitoring plan belong to the section owning inappropriate antidiuresis. The rate at which a low sodium may be corrected, and the demyelinating injury that follows correcting it too quickly, belong to the section owning primary polydipsia and mineral disturbance. Both are places where a confidently remembered figure is more dangerous than an admitted gap.

> **PITFALL**
>
> Treating the number instead of the patient, in either direction. A frightened team corrects a low sodium quickly because the value looks alarming, and causes an injury worse than the illness. Elsewhere a chronically low but stable sodium in a well patient is chased with intervention it never needed. Clinical state, tempo and neurological signs decide urgency; the value alone does not.

27.7 Managing the vignette: LOCUS, FOCUS and MODUS under exam conditions

Management answers here fail predictably: a competent psychiatric plan written for a patient whose problem is not, at this moment, psychiatric. LOCUS asks where care should happen and is answered first, because it governs everything after. Most of this work is outpatient and shared with a physician, but several presentations are emergencies in which a psychiatric bed cannot deliver the monitoring the body requires, and saying so explicitly earns the mark.

FOCUS asks what is treated and in what order. Acutely it is the derangement threatening the patient, with only as much behavioural containment as makes treatment possible. In the short term it is the underlying disease and the syndrome it drove. In the medium term it is whatever psychiatric picture remains after correction, where independent psychiatric treatment becomes justified rather than assumed. In the long term it is surveillance for recurrence and the accumulating physical cost of psychiatric treatment itself. MODUS then runs across four modes: pharmacological, carrying both the systemic treatment and a psychotropic chosen for the failing

organ; psychological, carrying adherence, dietary, activity and substance-use work; procedural, carrying surgical, replacement and dialytic intervention; and social, carrying cost, supervision, transport and the exposure that may need to change before the illness will.

- State the locus of care and the threshold at which this patient leaves the psychiatric setting.
- Name the immediate risk, often the derangement rather than the behaviour that prompted referral.
- Separate treatment of the disease from treatment of the syndrome, and say who owns each.
- Choose psychotropics for the organ that must clear them, and say what you would avoid.
- Define what will be monitored, by whom, and what will be done with the result.

IN INDIA

Two structural realities change these answers. The psychiatrist in a general-hospital unit is frequently the first and sometimes the only doctor to examine the patient physically, and onward referral is a plan rather than a guarantee, so the answer must contain a psychiatric management safe in the interval before the physician is actually reached. Patients also present late, so the fully developed somatic signature is more often available at first contact than the early subtle one: the pivot is easier to find and the treatment window more often closed. Long-term correction, whether replacement, chelation or a metabolic agent, is usually paid out of pocket, so a plan that ignores who will fund and supervise it for years is not a plan.

27.8 Writing the answer: the spine that survives exam pressure

Under pressure the reasoning collapses in the same order every time. The candidate names the psychiatric syndrome, recognises a disease, writes everything known about it, and never returns to the patient in the stem. The spine below is a retrieval aid rather than a diagnostic criterion, and asserts nothing this chapter has not already taught.

MNEMONIC

PIVOT

Pivot: the finding or cluster that most changes the ranking, and why. **I**mpostor: the psychiatric diagnosis this case would otherwise have received, and the systemic disease impersonating it. **V**erify: the investigation that separates them, and what each result would mean. **O**rder of correction: what is treated first, by whom, and how quickly it may safely be corrected. **T**hen reassess: what picture is expected to remain, when it will be reviewed, and what would reopen the formulation.

Be equally disciplined about what is missing. Say what was not in the stem and how it would change the ranking, since examiners write incomplete stems deliberately. Do not manufacture normal findings, and do not read silence in a short stem as evidence that a symptom is absent. Where the answer needs a threshold, rate, dose or interval that cannot be recalled with confidence, name the parameter, name the source that governs it, and move on. That is stronger

than a confidently wrong number, and better practice, because the number will be looked up before it reaches the patient.

Finish with what changes. A formulation that alters nothing is unfinished, and here something almost always changes: a test ordered today rather than next month, a drug stopped, a deficiency replaced before a window closes, a correction deliberately slowed, a transfer to a setting that can monitor the body.

> In every endocrine or metabolic vignette, name the syndrome, name the system that could have produced it, name the pivot that separates them, name the test that would settle it, and name the one action that cannot safely wait until the diagnosis is certain.

28 · Consolidation and quick review

Nothing in this material arrives labelled with the organ that caused it. These diseases have to be learned one system at a time, but that is not how they present. A patient arrives with low mood, a first psychotic episode, panic, apathy, a fall in cognitive performance or a change in behaviour the family cannot explain, and the organ responsible sits behind that presentation, silent or attributed to the mental illness. Consolidation means turning the material around: not from gland to symptom, which is how it was written, but from symptom to gland, which is how it is used.

Two further reversals matter as much. Part of this material is not bodily disease reaching psychiatry but bodily disease that psychiatry causes, where the exposure is a prescription written in good faith. And much of the benefit lies past the diagnosis, in what is corrected, how fast, what improves and who follows it up. Criteria, thresholds, doses and intervals belong to the sections that own them.

28.1 Read the material backwards, from syndrome to system

The psychiatric syndromes here are **final common pathways**. The brain has a limited repertoire of ways to fail, so disturbances as different as a depleted vitamin store, an excess of circulating steroid, a fallen sodium, an accumulated metal and a failing liver converge on the same handful of presentations: altered mood, drive, perception, arousal and cognition. The phenotype therefore under-determines the cause, which is why the mental state alone will never name the organ.

The usable consequence is an inverse index: for each presentation, the short list of systems capable of generating it, and the question that sorts that list fastest. It is recallable under pressure and more useful than a catalogue of glands, which is indexed by the one thing the patient does not bring you.

PRESENTATION	SYSTEMS CAPABLE OF GENERATING IT	THE QUESTION THAT SORTS IT FASTEST
Depressive syndrome	Thyroid disease; cortisol excess, endogenous or prescribed; adrenal insufficiency; vitamin B12 and folate deficiency; raised calcium; diabetes; organ failure	Weight direction, temperature tolerance, pigmentation, postural symptoms, the whole drug list
Mania or irritable overactivity	Prescribed corticosteroids; thyroid overactivity; thyroid hormone given for an affective indication	What was started, stopped or increased before the change
Psychosis	Copper disorder in a young patient; acute porphyria; thyroid disease; cortisol excess; vitamin B12 deficiency; the adult-onset inherited metabolic disorders; toxic exposure	What company it keeps: neurological signs, liver disease, abdominal crises, skin change, family history
Anxiety and paroxysmal autonomic symptoms	Thyroid overactivity; catecholamine-secreting tumour; hypoglycaemia; low calcium; hypoxia and carbon dioxide retention	Whether attacks are truly paroxysmal with observed autonomic change, and whether they track meals or exertion
Apathy and blunting	Thyroid underactivity; adrenal insufficiency; raised calcium; organ-failure encephalopathy; sustained toxic exposure	Whether the family describes slowing rather than sadness
Cognitive impairment	Vitamin B12 deficiency; thyroid underactivity; acute thiamine deficiency; hepatic and uraemic encephalopathy; copper disorder; recurrent hypoglycaemia; metal and solvent exposure	The tempo of decline, and whether anything correctable is still correctable
Acute confusion, stupor or catatonia	Severe derangement of sodium, glucose, calcium, oxygen, ammonia or urea; endocrine crisis; organ-failure encephalopathy; drug and toxin states	Which single derangement could account for the whole picture, and whether conscious level itself fluctuates
Behavioural change in a young adult	Copper disorder; the adult-onset inherited metabolic disorders; chronic occupational and environmental exposure	Family history, occupation, water source, home remedies, when function first slipped

RULE **The mental state names the syndrome, the body names the system, and only then does a test name the disease.** An answer jumping from a psychiatric presentation straight to a named endocrine or metabolic disease has skipped the step where the marks are, and clinically it is guesswork with a laboratory form.

28.2 The phenocopies this material exists to separate

A **phenocopy** is a systemic disease reproducing the surface of a psychiatric syndrome closely enough that the psychiatric diagnosis is made, recorded and then defended. These pairs are where examiners concentrate, because the discriminating evidence is usually available at first contact and simply not collected. Four discriminators do most of the work: tempo and periodicity; the company the symptom keeps, meaning observed physical signs rather than reported ones; its relationship to a meal, a fast, exertion, a menstrual phase or a new drug; and its behaviour under adequate psychiatric treatment.

THE PSYCHIATRIC LABEL USUALLY APPLIED	THE DISEASE IMITATING IT	WHAT SEPARATES THEM AT THE BEDSIDE
Panic disorder	Catecholamine-secreting tumour	Genuinely paroxysmal attacks with observed pallor, sweating and pressure change, without situational cueing or avoidance
Generalised anxiety in a young adult	Thyroid overactivity	Heat intolerance, weight loss with preserved appetite, resting tachycardia, tremor, eye or neck signs
Episodic anxiety, confusion or aggression	Hypoglycaemia	Timing to food, fasting or exertion, a glucose-lowering prescription, resolution with carbohydrate
Anxiety with paraesthesiae	Low calcium	Neuromuscular irritability on examination, with a precipitant such as neck surgery
Treatment-resistant depression in later life	Thyroid underactivity, vitamin B12 deficiency, raised calcium	Apathy and slowing out of proportion to reported sadness, with physical signs
A functional neurological presentation	Copper disorder	An abnormal movement of consistent form, liver disease, a specific eye sign, a family history
Relapse on maintenance treatment	Drug-induced hyponatraemia, organ-failure encephalopathy, hypoglycaemia	Impaired attention and fluctuating conscious level, which relapse rarely produces

TRAP **Explaining new physical or cognitive symptoms by the psychiatric diagnosis the patient already carries.** This is **diagnostic overshadowing**, the most reliable way this material is failed in examination and clinic alike. The longer and more severe the psychiatric history, the more readily new symptoms are absorbed into it, and that population carries the highest burden of the diseases described here.

28.3 Tempo sorts the differential; the window sets the urgency

Tempo is the most efficient discriminator available, and it costs nothing. Paroxysmal symptoms lasting minutes with objective autonomic accompaniment point to a secreting lesion or a fall in glucose or calcium. Deterioration over hours to a few days points to a solute, a glucose, an oxygen,

an ammonia or an endocrine crisis. Change over weeks to months fits a depleting nutrient store, an accumulating hormone excess or its withdrawal. Change over years fits an accumulating metal, a slowly failing organ or a sustained occupational exposure. Symptoms recurring with an external periodicity direct attention to whatever recurs around them: a menstrual phase, a dietary pattern or an intermittent exposure at home.

Tempo also sets urgency, because these diseases are correctable in four different senses and conflating them causes harm. The first must be corrected now or function is lost, and in some of them treatment properly precedes biochemical confirmation. The second must be corrected at a controlled rate, because the correction itself injures; that rate belongs to the section owning it and is never estimated at the bedside. The third is corrected and then waited on, because the syndrome lags the biochemistry, and judging response too early drives needless escalation of psychotropic treatment. The fourth is corrected expecting a remnant, so that a **residual syndrome** is recognised as a sequel to be rehabilitated rather than pursued as fresh pathology.

Inverse	Tempo	Class	After
SYNDROME TO SYSTEM	THE FASTEST SORTER	WHAT A TEST SETTLES	WHAT CORRECTION LEAVES

28.4 Know what each class of test can settle

Investigation here is not one activity but several, and marks are lost by treating every result as the same kind of evidence. A static level measures a circulating analyte at one instant and is shaped by biological rhythm, binding proteins, posture, fasting, the acute illness the patient currently has and the drugs already prescribed; it answers what the concentration was, not whether the system is regulated. **Dynamic testing**, by suppression or stimulation, answers the regulatory question directly, which is why it rather than the static level separates physiological adaptation from disease. A **functional marker**, a metabolite accumulating when a pathway is under-supplied, answers whether the tissue is genuinely deficient when the circulating level is equivocal. A timed collection answers output across an interval, so its validity depends on the collection being complete. Imaging answers structure and rarely establishes causation alone. Tissue and genetic testing answer cause, and reframe the family as well.

Two interpretive habits protect all of this. A **reference interval** describes a population and is not itself a decision threshold for the patient in front of you; the thresholds, the direction of change that matters and the situations in which the usual interval does not apply belong to the sections owning each analyte. And the pre-test question governs the reading, since the same borderline value means less in a patient screened without a hypothesis than in one whose history already pointed at that system.

PITFALL

The correct test, sent under conditions that make it uninterpretable. Several investigations here depend on the timing of the sample, the patient's clinical state, the completeness of a timed collection or how the specimen is protected in transport. Get that wrong and the result returns neither normal nor abnormal but useless, the patient has travelled home, and the working diagnosis hardens in the absence of evidence. Confirm the sampling conditions before the patient leaves.

28.5 Your own prescription is one of the exposures

Reverse the direction of causation and the psychiatrist stops being the detective and becomes the exposure. Psychotropic treatment moves weight, waist, lipids, glucose, prolactin and sodium, and it is prescribed to a population already carrying an excess burden of cardiometabolic disease and a shortened life expectancy. That is one of psychiatry's principal harms, and it is foreseeable at the moment of prescribing. Four questions belong before the pen moves.

- Which organ clears this drug, and is that the organ currently failing?
- What does the drug do to weight, glucose, lipids, prolactin and sodium, and what is this patient already carrying?
- What does the systemic disease do to the drug's margin of safety, through altered clearance, fluid shifts, seizure threshold, cardiac conduction or reduced marrow and hepatic reserve?
- What monitoring does this prescription commit me to, and can this patient complete it?

When an abnormal parameter appears, the decision turns on **attribution**. The candidates are the drug, the illness and its lifestyle, the pre-treatment baseline, and an unrelated intercurrent disease. Which one you choose determines whether you switch, add, reduce, treat or accept, and both reflexive answers are wrong: switching on a single value can cost a stable patient a hard-won remission, while accepting a value without recording that you accepted it leaves an untreated physical illness and no plan. The comparative liability of individual agents, the monitoring cadence, the criteria defining the metabolic syndrome and the management of established disturbance and raised prolactin are each stated once, in the owning section. Drug selection and receptor pharmacology belong to the Schizophrenia: management, antipsychotics and adverse effects notes, and lithium as a drug to the Mood disorders: pharmacotherapy notes.

GOOD TO REMEMBER

One sentence in the record does most of this work: which parameter changed, what its baseline was, what you attributed the change to, what you did, and when it will next be reviewed. A patient with no documented baseline cannot be shown to have deteriorated, and an accepted abnormality never written down is indistinguishable from one never noticed.

28.6 Management: one spine, whichever disease it is

Management questions here are almost always about a patient with both a systemic disease and a psychiatric syndrome, and they become answerable on the standard spine. **LOCUS** asks where

care happens. Most of this work is outpatient and shared, but a definable group of presentations are medical emergencies in which the derangement, not the psychiatric syndrome, dictates care on a medical ward or in critical care with psychiatric input rather than the reverse. Settling LOCUS early prevents the commonest structural failure here, admitting a physically deranged patient to a psychiatric bed that cannot deliver the correction the body needs. LOCUS is dynamic, and is revisited whenever the derangement, the risk or the supervision changes.

FOCUS asks what is treated, in what order. Acutely the target is the derangement that threatens the patient, with enough containment of behaviour and risk to make correction possible. In the short term it is the underlying disease and the syndrome it drove, expecting the second to follow the first. In the medium term it is whatever picture persists after correction, where independent psychiatric treatment becomes justified rather than assumed. In the long term it is relapse prevention for both illnesses and the cumulative cost of treatment itself.

MODUS asks which modes are in use: the systemic treatment and the psychotropic chosen for the organ that must clear it; adherence support, dietary and activity intervention, substance work and psychoeducation the family can act on; the procedural and replacement interventions the disease sections describe; and the social mode of cost, distance, supervision and the exposure that may have to change before the illness will. Running through all of it is **shared care**, which works only when someone is named as responsible for each parameter, the next review and the decision to escalate. Dose, eligibility and interval come from current disease-specific guidance, never from memory.

28.7 Follow-through: the part most answers omit

Correcting the abnormality is the middle of the story, and the reassessment that follows is where a formulation is confirmed or overturned. Four outcomes should be anticipated. The syndrome remits fully, so the systemic disease was the diagnosis and continued psychotropic treatment needs a reason. It improves partially, so the contribution was real but shared, and what remains is treated on its own terms. It does not change, so the two conditions were co-occurring and the psychiatric illness stands independently, which is a legitimate finding rather than a failed workup. Or the patient deteriorates during correction, raising the correction itself, an unmasked second process or a complication, and usually needing the physician back the same day.

Long-term follow-up then carries three jobs easily collapsed into one. It watches for recurrence of the systemic disease, which in several of these conditions announces itself first as a change in mental state and is misread as psychiatric relapse. It watches the parameters psychiatric treatment itself created, on the cadence set by the section owning the monitoring. And it revisits the explanation given to patient and family, since an account offered during an acute derangement is rarely the one they need a year later. Adherence deserves the scrutiny given to efficacy, because correction here means daily treatment and repeated travel for someone who feels well.

IN INDIA

The shared-care pathway that Western descriptions assume is frequently unavailable, and the psychiatrist may be the only doctor the patient sees at predictable intervals. The physical review, the repeat investigations and the coordination then fall to psychiatry, so they belong in the psychiatric plan rather than a referral letter whose outcome is never seen. Cost drives the rest: long-term replacement, chelation, glucose-lowering treatment and repeated monitoring are largely paid out of pocket, and a regimen the family cannot sustain is stopped silently rather than renegotiated, so ask what is affordable and stage the plan accordingly. Anthropometric thresholds for Indian patients differ from imported ones and are stated in the monitoring section. Since records are usually held by the patient rather than an institution, send the family home with a written summary of the correction, the parameters being followed and the signs that should bring them back.

28.8 The retrieval frame, and how to deploy it

When a stem is crowded with physical detail, frame the case before trying to solve it. The scaffold below is a retrieval aid, not a diagnostic criterion. Under pressure the steps most often dropped are the phenocopy question, because the label is already written, and the escalation question, because the physician is elsewhere.

MNEMONIC

PROBE

Phenocopy: ask which systemic disease could be imitating this syndrome before accepting the psychiatric label, especially in a patient who already carries one. **R**ate and rhythm: use tempo and periodicity to narrow the system, then ask which class of treatment window is open. **O**rgan: name the axis, nutrient, solute, metal, organ of clearance or exposure accounting for the whole picture rather than one symptom. **B**ind the test to the question: choose the investigation that answers what you asked, get the sampling conditions right, and read the result against this patient. **E**xposure and escalation: your prescription is an exposure, so decide what will be monitored, who is responsible, and when this patient needs a physician rather than a psychiatric review.

Keep the boundaries visible while doing it. Delirium as a syndrome, and the persistent amnestic syndrome that may follow acute thiamine deficiency, belong to the Stroke, TBI, delirium and focal brain syndromes notes. The dementia syndrome and its workup belong to the Dementias and neurocognitive disorders notes. Administration and scoring of the cognitive instruments belong to the Psychotherapies, clinical assessment, MSE and nosology notes, so cognitive change is described here in clinical language. Hormone synthesis, receptor signalling and neuroendocrine measurement belong to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes. Naming the interface and moving on shows the judgement being tested; re-teaching it shows the opposite.

Work backwards from the psychiatric syndrome to the systems that produce it, let tempo narrow the list and set the urgency, choose the test that answers your question, decide attribution before changing treatment, and follow the patient long enough to learn which of the two illnesses you were treating.

PLACEHOLDER references. Authored after grounding from the actual sources _CLAIMS.json rests on.

PART V

V

Stroke, TBI, Delirium and Focal Brain Syndromes

What a damaged brain does to a person, and how the syndromes of stroke, injury and delirium separate on the ward.

01 The cerebrovascular and post-stroke syndromes

02 Traumatic brain injury and its psychiatric sequelae

03 Delirium and the amnestic syndromes

04 The bedside examination and the focal lobe syndromes

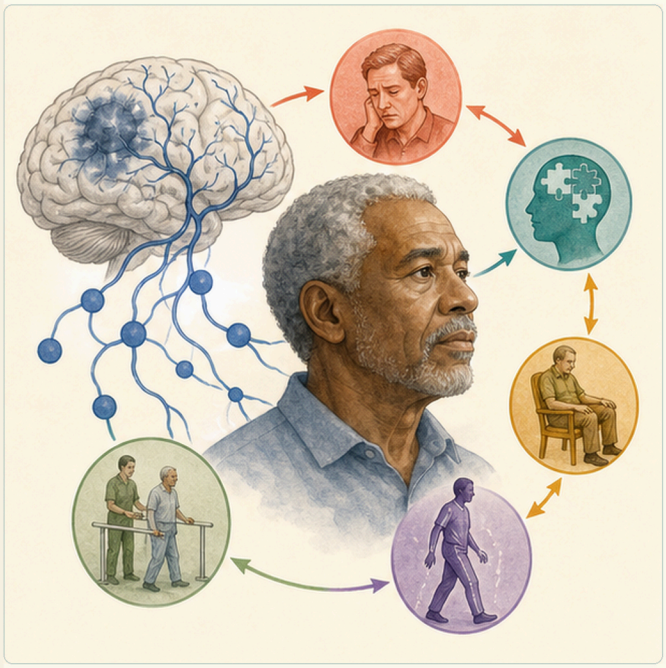
05 Structural lesions and the seizure overlap

06 Consolidate and apply

07 Sources and further reading

VISUAL 5.1

Stroke produces a network of psychiatric outcomes



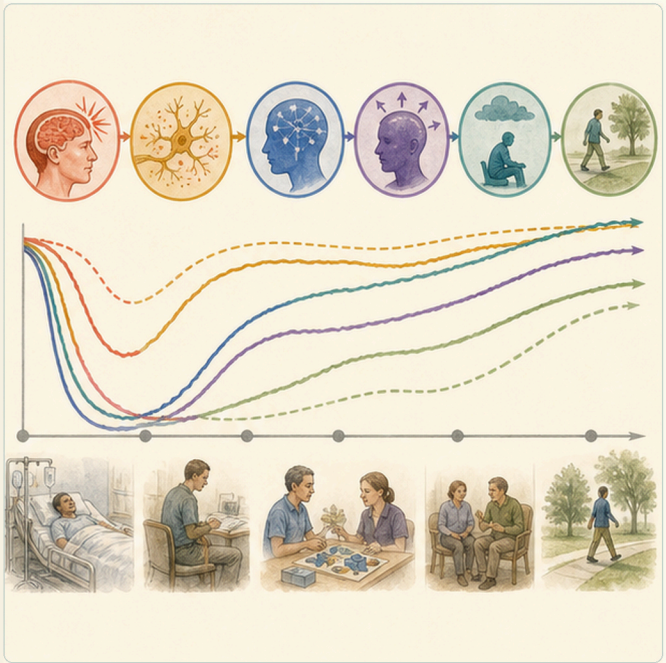
- Lesion and vascular burden
- Mood
- Cognition
- Behaviour and apathy
- Focal signs
- Rehabilitation

Post-stroke psychiatric outcomes reflect lesion, network disruption, disability, prior vulnerability and recovery context.

Source lock: day-34.final.html#sec-t06

VISUAL 5.2

TBI unfolds from injury into recovery



- Primary injury
- Secondary injury
- Cognition
- Behaviour
- Mood and trauma
- Rehabilitation

Traumatic brain injury begins with primary injury and evolves through secondary processes, neuropsychiatric sequelae and rehabilitation.

Source lock: day-34.final.html#sec-t07

VISUAL 5.3

Delirium runs on two simultaneous clocks



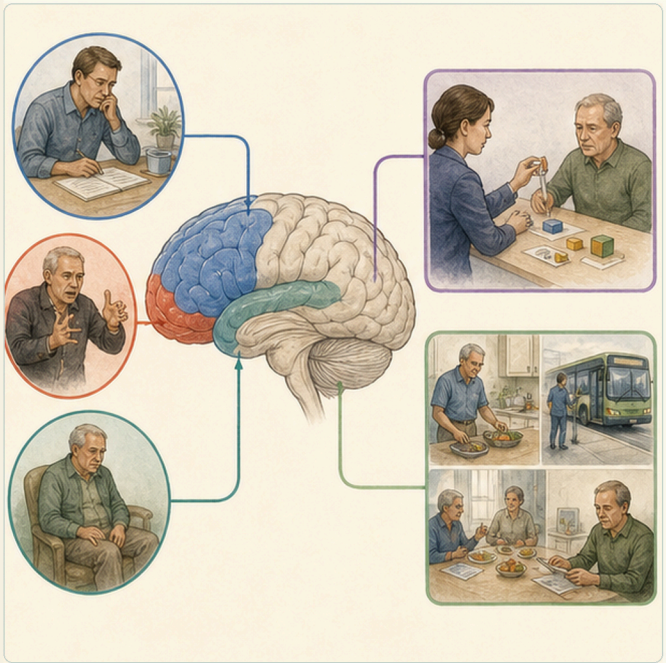
- Acute fluctuation
- Detect
- Find and treat causes
- Non-pharmacological care
- Medication restraint
- Prevention

Delirium care treats the underlying cause while supportive environmental measures, risk management and prevention run in parallel.

Source lock: day-34.final.html#sec-t16

VISUAL 5.4

Three frontal syndromes, three functional failures



- Dorsolateral
- Orbitofrontal
- Medial
- Executive testing
- Everyday function

Frontal syndromes are patterns of executive, inhibitory and motivational failure that require bedside testing and functional corroboration.

Source lock: day-34.final.html#sec-t21

PAPER IV

Stroke, TBI, delirium and focal brain syndromes

Five bands . one complete notes document

The cerebrovascular and post-stroke syndromes

1 · Acquired brain injury and focal brain syndromes for the psychiatrist: what this chapter owns, and what it points to

The psychiatric presentation may be the first visible expression of acquired brain dysfunction. A patient may reach the psychiatrist because of altered mood, belief, behaviour, memory, motivation, personality or awareness, while the decisive history is a vascular event, head injury, fluctuating medical state, focal cortical disorder, structural lesion or seizure-related change. The mental state describes the presentation. It does not, by itself, identify the neurological process that produced it.

This orientation provides the **clinical architecture** for the chapter. The dedicated sections teach each syndrome, lesion pattern, criterion set, investigation and treatment in depth. Here the task is to learn the shared method: establish the previous baseline, define the tempo and trigger, identify immediate danger, describe the dominant syndrome, localise only when findings converge, rank causes and consequences, and plan care across medical treatment, psychiatry and rehabilitation. This is where the marks sit, because a long list of associations is not yet a neurological psychiatric formulation.

1.1 Define the territory before naming the disorder

Acquired neurological psychiatry concerns changes in mental state, cognition and behaviour that arise in the context of brain injury, diffuse cerebral dysfunction or structural neurological disease. Its central difficulty is overlap. Affective symptoms, psychosis, apathy, disinhibition, amnesia and functional decline can each occur through several mechanisms, and the same neurological illness can produce different psychiatric phenotypes in different people.

The opening question is therefore not “Which psychiatric diagnosis is this?” It is “What changed, from what baseline, in what neurological and medical context?” This creates an **index problem** that is clinically useful without claiming more than the evidence supports. “New behavioural and cognitive change after a neurological event” preserves the sequence and uncertainty. A premature label such as depression, dementia or psychosis may discard the very observations needed to distinguish neurological injury, systemic disturbance, treatment effect, adaptation and independent psychiatric illness.

The chapter crosses several levels of explanation. At the symptom level, it describes what the patient experiences and what others observe. At the syndrome level, it asks which features cohere. At the anatomical level, it considers diffuse dysfunction, focal injury, disconnection and network disturbance. At the aetiological level, it asks what process caused the state. At the

functional level, it measures the effect on safety, autonomy, rehabilitation and relationships. A defensible answer moves between these levels without confusing them.

This approach also prevents diagnostic overshadowing in either direction. A known psychiatric disorder should not make a new neurological change invisible. Conversely, a visible lesion should not be made responsible for every symptom merely because it is concrete. Pain, sleep disruption, medication, sensory loss, fear, grief, family strain and premorbid vulnerability remain clinically active even when neurological disease is established.

RULE **Describe the change before naming the syndrome, and describe the syndrome before assigning a site or cause.** Each later inference should explain the observations already established rather than replacing them.

1.2 Read the chapter as a set of clinical lenses

The chapter is arranged by the kind of neurological problem confronting the psychiatrist. These are not sealed compartments. A vascular event may produce a focal syndrome, an injured patient may become delirious, a structural lesion may alter cognition, and seizures may complicate several forms of brain injury. The value of the map is to show which lens leads and where detailed teaching is located.

CLINICAL TERRITORY	OPENING QUESTION	DEPTH SUPPLIED BY THE CHAPTER
Cerebrovascular and post-stroke syndromes	What changed after the vascular event, and how securely can it be attributed?	Affective, motivational, psychotic, expressive, awareness-related and vascular cognitive presentations
Traumatic brain injury	How do injury severity, evolving damage and recovery context shape the psychiatric course?	Classification, concussion, psychiatric and personality sequelae, cognition, rehabilitation and pharmacological principles
Delirium and amnestic syndromes	Is the central disturbance an acute unstable state, a persistent memory syndrome or a transient amnesia?	Recognition, mechanisms, investigation, management, prevention, amnestic pathology and confabulation
Bedside and focal brain syndromes	Do cognitive and behavioural findings converge on a cortical, subcortical or disconnection pattern?	Ordered bedside examination and the frontal, temporal, parietal, occipital and disconnection syndromes
Structural lesions and seizure overlap	Could mass effect, altered pressure, an extra-axial collection, tumour, hydrocephalus or seizure-related illness explain the presentation?	Reversible structural differentials, location-related presentations, pressure syndromes and psychiatric implications of seizures

Event	**State**	**Network**	**Impact**
WHAT HAPPENED	HOW THE BRAIN FUNCTIONS NOW	WHERE FINDINGS CONVERGE	WHAT CARE MUST CHANGE

Use the map from either starting point. A known neurological event directs the clinician toward likely psychiatric and cognitive sequelae. An unexplained psychiatric change directs the clinician back toward neurological possibilities that must not be missed. The sections supply depth, but the orientation preserves integration: event, current state, anatomical pattern, cause, function and trajectory belong in the same case formulation.

1.3 Let tempo and trajectory organise the differential

Tempo is the most efficient opening discriminator. Abrupt onset, fluctuation, a stepwise course, gradual progression, episodic recurrence and a stable residual deficit raise different clinical questions. The sequence of features matters as much as their presence. A behavioural change that follows altered alertness, a focal deficit, injury or a seizure-like event is framed differently from a longstanding psychiatric syndrome later accompanied by neurological illness.

Start with the last reliably known baseline. Establish prior cognition, personality, daily function, psychiatric history, substance exposure and treatment. Place neurological, psychiatric, cognitive and systemic changes on one timeline. Mark injuries, procedures, medication changes, intoxication or withdrawal, intercurrent illness, transfers between settings and changes noticed after return home. A family statement that the patient was “normal before” should be unpacked into actual capacities and roles.

Distinguish onset from recognition. Executive or memory difficulty may become visible only when environmental demands increase. A family may date all later problems to a dramatic injury even when some were present beforehand. Conversely, a polished interview may conceal a clear decline from a highly functioning baseline. The patient should be compared with their own previous level rather than with an abstract norm.

State dependence means that performance and behaviour can vary with arousal, pain, fatigue, anxiety, intoxication, withdrawal, medication, sensory impairment and task demands. A snapshot taken during acute illness should not be treated as a fixed estimate of future ability. Serial observation tests the formulation: improvement with medical stabilisation, persistence despite recovery, or emergence under complex demands each changes the causal balance.

The trajectory also guides urgency. Rapid deterioration or fluctuation demands a different response from a stable deficit. The formulation should state what is changing, what appears fixed, what may be reversible, and which uncertainty requires reassessment. Avoid precise prognostic language unless the relevant disease section supplies evidence for it.

1.4 Hold parallel formulations rather than one grand diagnosis

A robust case uses **parallel formulation**. The syndrome, anatomical, aetiological and consequence accounts are developed together, but each answers a different question. They constrain one another without being collapsed. When they do not align, the discrepancy is information, not an invitation to force certainty.

The **syndrome formulation** describes the dominant pattern in clinical language: altered arousal and attention, amnesia, executive-behavioural change, a focal higher-function deficit, affective or

psychotic symptoms, or a mixed presentation. It separates observed behaviour from inferred intention. Sparse speech, inactivity, emotional display or inaccurate recall should be described before being interpreted.

The **anatomical formulation** asks whether the pattern suggests diffuse cerebral dysfunction, a lateralised deficit, a focal cortical syndrome, disconnection, subcortical network involvement or pressure-related compromise. The hypothesis becomes stronger when several findings converge and weaker when it rests on an isolated test error or an imaging abnormality without clinical correspondence.

The **aetiological formulation** ranks direct injury, evolving complications, systemic illness, medication and substances, seizures, structural disease and independent psychiatric conditions by urgency, probability and reversibility. More than one cause may be active. A common explanation may lead the probability list while a less likely but dangerous alternative leads the action list.

The **consequence formulation** covers risk, decision-making, self-care, treatment adherence, caregiver burden, rehabilitation participation and social role. It prevents a diagnostically elegant answer from ignoring what must change today. It also exposes cases where a seemingly mild symptom has major functional effect, or a conspicuous symptom causes distress without equivalent loss of independence.

> **PITFALL**
>
> Inferring intention from impaired performance. “Refusal,” “lack of effort” or “uncooperativeness” may instead reflect poor comprehension, weakness, pain, fatigue, altered awareness, fear or inability to execute the requested act. Record what happened and what the task required before assigning motive.

1.5 Know what this chapter owns and what it hands over

Neurological psychiatry touches constructs that are taught elsewhere in full. A disciplined answer names the interface and cross-refers; it does not reproduce another chapter’s criteria, scoring table or pharmacology. This boundary is part of clinical competence because it keeps the answer centred on the neurological question actually asked.

- The ordered bedside cognitive examination is taught here, while standardised cognitive instruments, their item content and scoring belong to the Psychotherapies, clinical assessment, MSE and nosology notes. Use instruments without reproducing their proprietary or chapter-owned detail.
- Vascular and post-injury cognitive presentations are taught here through their neurological lens. The general dementia syndrome and its full diagnostic workup belong to the Dementias and neurocognitive disorders notes.
- Substance-withdrawal delirium is placed within the delirium differential here, while withdrawal timelines, severity scoring and regimen detail belong to the Substance use

disorders notes. Nutritional emergencies that precede persistent amnesia are treated in the Endocrine and metabolic psychiatry notes.

- The psychiatric implications of post-injury and post-stroke seizures are taught here. Seizure classification and anticonvulsant pharmacology belong to the Epilepsy and psychiatry notes, while antidepressant pharmacology and general drug selection belong to the Mood disorders: pharmacotherapy notes.

The same discipline applies within this chapter. Post-stroke affective syndromes, traumatic brain injury classification, delirium criteria, amnestic pathology, focal lobe syndromes, structural lesions and seizure-related presentations each have one teaching home. An orientation may name them and explain how to approach the territory, but their numbers, criteria, mechanisms and treatment evidence remain with the owning section.

A useful handover sentence does real work: name the neighbouring construct, identify why it matters to the current case, and point to its owning chapter. Then return to the neurological lens. This is shorter and more accurate than inserting a generic dementia workup, a copied cognitive test, a withdrawal regimen or a drug-selection essay into every answer that happens to touch them.

1.6 Localise cautiously and think in networks

Network localisation is a hypothesis built from converging observations. Begin with the descriptive phenotype and neurological findings, then ask whether they form a coherent pattern. A localising account should explain the combination rather than match one dramatic feature. An isolated bedside error is insufficient because performance may be distorted by demands outside the ability being examined.

Structural imaging contributes evidence about lesion location, vascular burden, mass effect and other pathology, but it does not replace bedside description. An abnormal scan may be causally important, contributory or incidental. Concordance among tempo, phenotype, examination and imaging is what makes attribution persuasive. When these disagree, preserve the discrepancy and reconsider the formulation.

Psychiatric syndromes rarely map onto one anatomical point in a deterministic way. Focal injury can disturb connected regions, diffuse injury can produce apparently focal deficits, and similar behaviours can arise through different network failures. Premorbid organisation, reserve, sensory and motor disability, medication and environment shape expression. This is why lesion associations should generate and test hypotheses rather than end discussion.

TRAP **Do not turn an association into a localisation law.** A remembered side, lobe or structure may support an answer, but it cannot substitute for phenomenology, tempo, neurological examination and competing explanations.

The bedside cognitive method, executive testing and focal higher-function examination are taught in their dedicated sections. The orientation contributes only the interpretive rule: before calling a deficit localising, confirm that the observed performance is valid for the ability being

inferred. When validity is uncertain, seek converging observations rather than forcing a site from the task.

1.7 Assess safety and manage across the recovery pathway

Assessment begins with immediate safety rather than diagnostic elegance. Physiological instability, reduced or fluctuating consciousness, rapid deterioration, a new focal change, seizure-like events, recent injury or other evidence of an evolving neurological or systemic emergency should move care toward urgent medical evaluation. These are escalation signals, not diagnoses. Seek collateral information when the neurological or cognitive state may make the history unreliable.

Management follows the shared spine. **LOCUS** asks where care can safely occur: outpatient and rehabilitation settings for stable patients with adequate support, inpatient care when medical or behavioural needs exceed community capacity, and emergency or higher-dependency care when instability or deterioration demands it. The setting should match the most urgent active problem, not the specialty of the clinician who first received the referral.

FOCUS orders targets over time. Acute work addresses stabilisation, the underlying cause, immediate risk and behaviour that prevents essential treatment. Shorter-term work supports sleep, pain control, communication, mobilisation and participation in rehabilitation. Continuing care addresses residual cognitive, affective and behavioural problems, recurrence prevention, social reintegration, occupational goals and caregiver strain. Targets should be observable and linked to function.

MODUS ensures that care is not reduced to prescribing. Cause-directed medical treatment, cautious symptom treatment, psychological support, rehabilitation, environmental adaptation, nursing strategies, family work and social intervention may all be required. Procedural treatment belongs where the underlying neurological condition indicates it. Current guidance and local protocols should direct disease-specific decisions; no named guideline is available to this orientation.

> **GOOD TO REMEMBER**
>
> Ask what the patient can do across a normal day, not only what can be demonstrated in a quiet interview. Nursing, therapy and caregiver observations often reveal participation, supervision needs and risk more clearly than a symptom label does.

Monitoring should cover benefit and harm. Define the target behaviour or function, identify who can observe it, and state what worsening prompts reassessment. Avoid unnecessary sedation and other interventions that obscure examination or reduce rehabilitation participation. Family involvement is essential, but clinical responsibility should not be transferred to caregivers without support, explanation and a workable follow-up route.

1.8 Use the NEURAL scaffold in essays and cases

Under examination pressure, preserve the order of reasoning. The master mnemonic is **NEURAL**. It is an educational scaffold rather than a diagnostic instrument. It can organise a vascular case, head injury, delirium, amnestic presentation, focal syndrome, structural lesion or seizure-related psychiatric problem without replacing the disease-specific knowledge in the relevant section.

MNEMONIC

NEURAL

New change from baseline: describe what altered and what was present before. **E**volution and event: reconstruct tempo, trigger, sequence and trajectory. **U**rgency: identify instability, deterioration and immediate risk. **R**egional or network hypothesis: localise only when findings converge. **A**etiology and alternatives: rank direct injury, complications, systemic causes, treatment effects and independent illness. **L**ife impact and locus: connect function, rehabilitation, setting, treatment modes and follow-up.

For a short case, compress NEURAL into a problem representation, leading formulation, dangerous alternative, discriminating findings and immediate plan. For a long essay, expand each element and insert disease-specific epidemiology, criteria, investigations and treatment from the owning section. In a viva, make calibration audible: state what is observed, what is inferred, what remains uncertain and which finding would change the ranking.

The enduring habit is integration without overreach. Psychiatric symptoms can be neurologically meaningful without being anatomically unique. A lesion can be clinically important without explaining the whole patient. A fluent conversation can coexist with substantial cognitive or functional impairment. A dramatic behaviour can be less urgent than quiet deterioration. A good answer keeps these tensions visible and resolves them through collateral history, ordered examination, appropriate investigation and serial review.

NEURAL: define the new change, reconstruct evolution and event, detect urgency, build a regional or network hypothesis, rank aetiology and alternatives, then plan for life impact and locus of care.

2 · Post-stroke depression: epidemiology, lesion correlates, diagnosis and treatment

Why it matters clinically. Post-stroke depression is not an understandable sadness that can be dismissed as an inevitable reaction to disability. It is a depressive syndrome occurring in the biological, psychological and social aftermath of stroke, with consequences for quality of life, disability and survival. The examiner therefore expects more than a prevalence figure: define the syndrome, show why detection is difficult, present lesion evidence without false certainty, and separate treatment of depression from attempts to improve neurological recovery.

IN INDIA

Screen for and treat post-stroke depression as a disorder in its own right, choosing an SSRI such as sertraline or escitalopram, and do not wait for neurological recovery before addressing the mood.

The accompanying figure situates depression among the wider post-stroke neuropsychiatric syndromes. Post-stroke mania, anxiety, apathy, fatigue, psychosis, emotional incontinence and denial of deficit are named here only because they enter the differential; their phenomenology and treatment belong in their dedicated sections. Vascular cognitive impairment is likewise a separate construct.

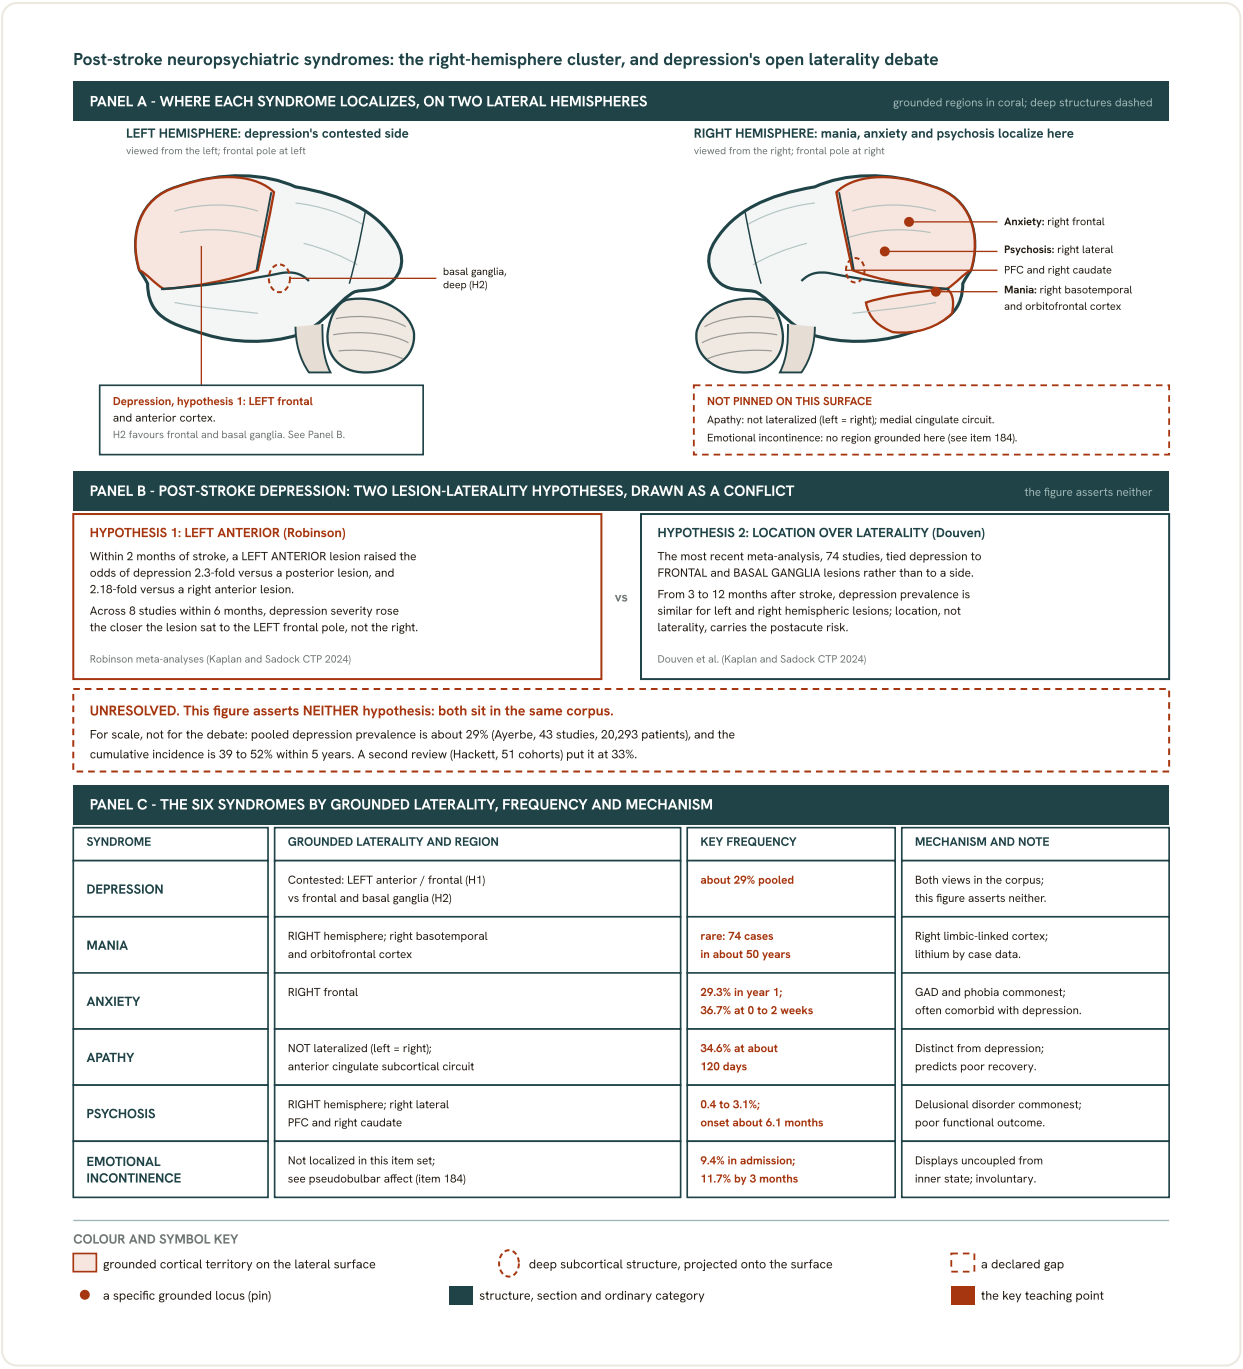


Fig 34.1 **The post-stroke neuropsychiatric map.** Mania, anxiety and psychosis cluster in the RIGHT hemisphere (right basotemporal and orbitofrontal for mania, right frontal for anxiety, right lateral prefrontal cortex and right caudate for psychosis); apathy is not lateralized and emotional incontinence has no region grounded in this item set. Depression is the one open laterality question, shown as a labelled conflict between the left-anterior hypothesis and the location-over-laterality view, with neither asserted. (Kaplan and Sadock Comprehensive Textbook of Psychiatry 2024; Tasman Psychiatry; Cognitive Assessment for Clinicians, Hodges.)

RULE **Diagnose a depressive syndrome, establish its relationship to stroke, and treat the mood disorder as a clinically important outcome in its own right.** Do not require a particular lesion side, and do not promise neurological recovery merely because depression responds.

2.1 Epidemiology: common, persistent and method-dependent

Post-stroke depression is common across the stroke pathway rather than confined to one ward or one phase. In the review by Ayerbe and colleagues, pooled prevalence at any assessed time was **29%**; prevalence did not differ significantly across time points or settings in that synthesis.

Cumulative incidence reached 39% to 52% within 5 years, showing that a patient who is well during initial rehabilitation remains at later risk.

29%	33%	39% to 52%	5 years
POOLED PREVALENCE	SECOND LARGE REVIEW	CUMULATIVE INCIDENCE	INCIDENCE WINDOW

A second large review by Hackett and colleagues estimated depression in 33% of stroke patients. The closeness of the pooled estimates is useful, but neither should be treated as a universal ward rate. The denominator, care setting, assessment method and inclusion of syndromal versus subsyndromal depression change what is counted. Prevalence is therefore best presented as a range of evidence, followed by the reasons for heterogeneity.

CARE SETTING	MAJOR DEPRESSION	MINOR DEPRESSION	INTERPRETIVE POINT
Community-based samples	14.1%	9.1%	Community estimates need not resemble hospital case-mix.
Acute or rehabilitation hospitals	21.6%	20.0%	Active medical and rehabilitation settings contain substantial syndromal and subsyndromal morbidity.
Outpatient studies	24.0%	23.9%	Ongoing surveillance remains necessary after discharge.

These setting-specific estimates come from the systematic review by Robinson and coinvestigators. Their practical lesson is not that one setting is intrinsically more depressogenic, but that sampling and clinical context shape the observed rate. A good answer therefore states that depression is frequent, gives a defensible pooled estimate, and then explains heterogeneity rather than averaging unlike populations.

2.2 Time course, disability and mortality

Risk begins early but does not end with acute admission. Depression is particularly frequent during hospitalisation and the early post-stroke period, while observed rates may remain stable for as long as 10 years. This coexistence of early frequency and long persistence explains why a single negative screen is insufficient. Case finding must recur when function, dependence, living circumstances or engagement with rehabilitation changes.

The course is heterogeneous. In one longitudinal estimate, mean duration was 39.0 plus or minus 31.8 weeks for major depression and 12.2 plus or minus 18.2 weeks for minor depression. Another study found a shorter mean duration for major depression but a longer mean course for minor depression, showing that the labels do not determine an individual trajectory. Among patients depressed at 3 months, **50% had recovered by 1 year**; across a meta-analysis, recovery at that point ranged from 15% to 57%. Variation reflects more than symptom severity: ascertainment, recurrence, medical course and follow-up loss all influence the observed pattern.

Outcome associations are clinically weighty. Post-stroke depression is associated with poorer quality of life, greater disability and later mortality. In one long follow-up, major or minor depression was associated with a **3.4-fold increase in the odds of mortality** compared with no depression. In another cohort, depression assessed early after stroke was associated with 50% higher mortality at 1 year. These observational associations do not by themselves establish a single causal pathway, but they make therapeutic nihilism unsafe.

GOOD TO REMEMBER

Repeat assessment across transitions of care. Ask again when the patient returns home, rehabilitation stalls, dependence becomes clearer, or the caregiver reports withdrawal. A negative interview during acute admission does not close the diagnostic window.

2.3 Lesion correlates: the high-yield unresolved debate

The classic account linked depression to anterior left-hemisphere injury. In a meta-analysis examining depression within 2 months of stroke, a **left anterior lesion** increased the odds of depression by 2.3 compared with a left posterior lesion and by 2.18 compared with a right anterior lesion. A separate synthesis found that depression severity increased as lesions lay closer to the left frontal pole, without the same relation to the right frontal pole. This became the memorable left-anterior hypothesis.

Later synthesis makes a simple laterality rule difficult to defend. The most recent meta-analysis described in the available evidence included 74 studies and identified clearer risk with **frontal and basal ganglia lesions**. An associated imaging review emphasised lesion location rather than laterality or lesion type in the postacute phase. In studies conducted 3 to 12 months after stroke, prevalence was similar after left- and right-hemisphere lesions, although lesion location appeared more relevant during the earlier months.

The defensible synthesis is phase-sensitive and probabilistic. Frontal-subcortical systems plausibly connect affect regulation, motivation, cognition and behavioural initiation, so disruption may contribute to depressive expression. Yet stroke also alters disability, communication, identity, relationships and reinforcement. Lesion anatomy is therefore one component of a biopsychosocial formulation, not a bedside diagnostic test.

TRAP **Do not answer "post-stroke depression equals a left frontal lesion."** The classic left-anterior association and the later location-over-laterality synthesis are genuinely in tension. State both, preserve their time windows, and conclude that no single hemispheric rule is adequate.

Equally, absence of a frontal or basal ganglia lesion does not exclude depression. Imaging can support the neurological formulation and may help explain associated deficits, but the depressive diagnosis rests on phenomenology, course, impairment and causal judgement. Anatomical correlation should never replace direct enquiry about mood, pleasure, hopelessness and safety.

2.4 Phenomenology and diagnostic classification

Comparative studies suggest that the phenomenology of post-stroke depression resembles depression in the general population and that **the same diagnostic criteria** can be used. This is important because clinicians sometimes invent a reduced "organic depression" syndrome consisting only of tearfulness and poor rehabilitation. The assessment should instead seek the full depressive pattern, including mood change, loss of interest, cognitive and neurovegetative symptoms, and associated distress or impairment, while interpreting each symptom in its neurological context.

Formal current classification should lead with ICD-11; the verified DSM-5 formulation here is a cross-map rather than a replacement for primary ICD wording. DSM-5 places the syndrome among depressive disorders due to another medical condition. When the post-stroke syndrome meets full criteria for a major depressive episode, the formulation is **depressive disorder due to stroke, with major depressive-like episode**. Subsyndromal depression is classified as depressive disorder due to stroke, with depressive features; the description supplied here requires more than two but fewer than five relevant symptoms persisting for at least 2 weeks.

The crucial attribution method is **etiological symptom counting**: a symptom is counted only when it is not judged to be caused by the physical illness itself. This avoids both extremes. Counting every post-stroke sleep, appetite, energy or concentration change as depressive can overdiagnose; automatically discarding every somatic or cognitive symptom can underdiagnose. The clinician asks whether the symptom belongs to a coherent depressive syndrome, whether it changed with affective experience, and whether a neurological, medical or environmental explanation is better.

MNEMONIC

M-I-N-D

Mood and interest are assessed directly; **I**mpairment is described against the post-stroke baseline; **N**eurological overlap is judged symptom by symptom; **D**anger, including hopelessness, self-neglect and suicidal thinking, is assessed rather than inferred from appearance.

2.5 Bedside diagnosis and differential diagnosis

The interview must be adapted, not abandoned. Aphasia, dysarthria, slowed processing, sensory loss, motor disability and fatigue can reduce the amount or speed of information obtained. Use brief language, allow response time, establish a reliable communication mode, obtain collateral history and compare behaviour with the premorbid and immediate post-stroke baseline. Communication difficulty lowers confidence in a negative interview; it does not prove absence of depression.

The examination should integrate several evidence streams:

- Directly elicit persistent mood change, reduced interest or pleasure, hopelessness, guilt, worthlessness and thoughts of death or self-harm.

- Describe biological and cognitive symptoms, then judge whether each is better explained by the stroke, another medical problem, treatment burden or the depressive syndrome.
- Ask the family about withdrawal, loss of initiative, reduced emotional responsiveness, self-neglect and change from the patient's usual coping style.
- Assess capacity to participate in treatment, current supports, caregiver strain, adherence barriers and the effect of mood on rehabilitation goals.

Apathy, fatigue, emotional incontinence, delirium and vascular cognitive impairment can coexist with or be mistaken for depression. Name them in the differential, but assess each through its dedicated section rather than importing its phenomenology into the depressive formulation. The immediate task here is to decide whether a coherent depressive syndrome is present in addition to the neurological condition.

PITFALL

Do not diagnose or exclude depression from facial expression, crying or rehabilitation performance alone. Neurological impairment can alter expression and performance in either direction. The diagnosis requires triangulation across subjective experience, observed behaviour, collateral change, neurological findings and longitudinal course.

Safety assessment remains explicit even when communication is impaired. Explore suicidal thinking, intent, self-neglect, refusal of essential care and access to support using the most reliable communication method available. If the patient cannot participate adequately, document the limitation, increase observation and collateral assessment, and reassess. The mortality association strengthens the case for active detection but should not be misrepresented as proof that every death is directly caused by depression.

2.6 Management framework: locus, focus and modus

Locus. Most medically stable patients can be assessed and treated through coordinated outpatient stroke, rehabilitation and mental health care. Inpatient care is appropriate when stroke-related medical needs, severe depression, inability to maintain safety, profound self-neglect or inability to participate in essential care require close supervision. Emergency assessment is required for immediate danger, rapidly changing neurological status or an acute confusional presentation. The site of care should be revised as medical stability, communication and social support change.

Focus. In the acute phase, establish the syndrome, exclude urgent mimics, address safety and preserve engagement with stroke care. Over the short term, target depressive symptoms, sleep-wake routine, distress, treatment participation and caregiver understanding. In the middle phase, review response, adherence, residual disability, role loss and rehabilitation goals. Over the long term, monitor recurrence, persistent symptoms, social isolation and transitions in support. Mood outcome and neurological outcome should be documented separately.

Modus. Treatment is multimodal: psychoeducation, collaborative goal setting, psychological intervention, social and caregiver support, rehabilitation coordination and antidepressant

treatment when indicated. Communication adaptations and supported decision-making allow psychological work to continue despite disability. Problem-solving approaches can address practical losses and reduced reinforcement without implying that the depression is merely situational. Current stroke and prescribing guidance should be consulted for patient-specific safety, interactions and monitoring.

The post-stroke evidence supports treating clinically significant depression, particularly severe illness. It does not remove the need to individualise choice against age, medical comorbidity, concomitant treatment, swallowing, cognition, previous response and the patient's priorities. General antidepressant mechanisms, comparative selection and broader prescribing principles belong in the Mood disorders: pharmacotherapy notes; the present section owns their use and evidence after stroke.

2.7 Antidepressant and psychological treatment evidence

Controlled trials establish that depressive symptoms after stroke can respond to medication. The early nortriptyline trial found greater improvement in depression ratings with nortriptyline 50 to 100 mg/day for post-stroke depression than with placebo over 6 weeks. The first placebo-controlled SSRI trial likewise found greater improvement with citalopram, using citalopram 20 mg/day under age 65 and 10 mg/day over age 65 for post-stroke depression, assessed over 6 weeks. These are trial regimens for this indication, not universal prescriptions.

EVIDENCE QUESTION	FINDING	CLINICAL INTERPRETATION
Nortriptyline versus placebo	Greater reduction in depression ratings	Supports antidepressant efficacy in established post-stroke depression.
Citalopram versus placebo	Greater reduction in depression ratings	Provides controlled SSRI evidence for the mood indication.
Pharmacotherapy versus control	Meta-analysis favoured pharmacological treatment	Supports treatment, particularly for severe depression.
SSRIs in later meta-analysis	Mood benefit remained evident	Reinforces mood efficacy without proving neurological benefit.
Problem-solving therapy for prevention	Lower incidence than placebo	Psychological treatment has preventive evidence in a selected nondepressed acute-stroke sample.
Antidepressant treatment and survival	Higher survival at extended follow-up	Suggestive long-term association, not a reason to promise survival benefit to an individual.

The pharmacotherapy synthesis included 16 trials and 1,655 participants and favoured active treatment with an odds ratio of 0.47. A later meta-analysis including 5,370 patients reported a relative risk of 0.78 for patients taking SSRIs. These estimates support efficacy for depression, but response still requires repeated measurement of symptoms, function, tolerability and safety.

A follow-up of patients treated with nortriptyline or fluoxetine for 12 weeks found survival of 59.2% among actively treated patients and 36.4% among placebo recipients at extended follow-up. This finding is clinically interesting but should be expressed cautiously: treatment allocation, later health and comorbidity complicate causal interpretation. The immediate evidence-based reason to treat remains relief of depression and restoration of engagement, not a guaranteed mortality effect.

2.8 Prevention and the functional-recovery controversy

Prevention is distinct from treatment of established depression. In nondepressed acute-stroke patients, a trial used escitalopram 5 mg/day over age 65 and 10 mg/day under age 65 for prevention of post-stroke depression over 1 year. Depression developed in 8.5% of the escitalopram group, 11.9% of the problem-solving therapy group and 22.4% of the placebo group. The result shows preventive efficacy in that trial population, but it does not mean every stroke survivor should automatically receive medication. Baseline depression status, risk, preference, comorbidity and monitoring capacity still matter.

The more important examination distinction is between mood treatment and neurological recovery. Some earlier findings associated fluoxetine 20 mg/day for post-stroke recovery, given for 3 months or nortriptyline 100 mg/day for post-stroke recovery, given for 3 months with improvement assessed later, independent of depression. However, the later systematic review concluded that SSRIs do not appear to improve functional recovery from neurological impairment after stroke. These findings are discordant, so the safest conclusion is deliberately narrow.

Prescribe an antidepressant to treat or, in selected circumstances, prevent depression. Do not present it as a routine neurorestorative treatment. Rehabilitation continues according to neurological and functional need, while mood response is assessed on its own axis. If improvement in participation follows relief of depression, that is clinically valuable, but it is not equivalent to a direct drug effect on motor recovery.

A high-scoring synthesis therefore moves in a disciplined sequence: establish the substantial and persistent burden; state that ordinary depressive phenomenology remains valid; apply etiological judgement to overlapping symptoms; present the left-anterior hypothesis alongside later location-focused evidence; assess safety and mimics; and use coordinated psychological, social, rehabilitation and pharmacological care. Close by separating three questions that are often confused: treatment of established depression, prevention in a nondepressed patient, and enhancement of neurological recovery.

> Post-stroke depression is common, diagnostically clinical and treatable; lesion side is not a rule, and antidepressant mood efficacy must not be converted into a claim of neurological recovery.

3 · Post-stroke mania, anxiety, apathy and fatigue

Why this cluster matters. A survivor who is restless, fearful, inactive or exhausted may be showing a direct neuropsychiatric consequence of stroke, an understandable response to disability, a coexisting depressive syndrome, a medical complication, or several of these at once. The examination marks lie in separating syndromes that look superficially alike: excess drive in mania, threat anticipation in anxiety, reduced initiation in apathy, and reduced capacity in fatigue. The clinical marks lie in making the same distinction before choosing treatment.

> **IN INDIA**
>
> If secondary post-stroke mania needs pharmacotherapy, lithium may be used in selected patients; a specialist opinion and appropriate baseline monitoring are worthwhile before starting.

Post-stroke mania is striking but rare; **post-stroke anxiety** and **post-stroke apathy** are much more frequent. Fatigue is also a central clinical complaint, but a defensible prevalence or time-course estimate is not available in the evidence represented here. Do not fill that gap with a remembered figure. Across all four presentations, diagnosis requires a syndrome-level description, temporal relation to stroke, functional effect, collateral history and a deliberate search for competing explanations.

74	29.3%	34.6%	23
ADULT MANIA CASES REPORTED OVER ABOUT 50 YEARS	POOLED ANXIETY PREVALENCE IN THE FIRST YEAR	MEAN APATHY PREVALENCE AT A 120-DAY MEDIAN	PATIENTS IN THE LONGITUDINAL MANIA COURSE STUDY

3.1 A four-syndrome clinical map

Begin with the patient's dominant change, not with a diagnostic label. Mania is an activated state: mood, speech, thought, sleep and behaviour move together towards excess. Anxiety is organised around threat, uncertainty, bodily alarm or avoidance. Apathy is a disorder of spontaneous goal direction: the patient does less because initiation, interest or emotional engagement has diminished. Fatigue is experienced as excessive effort, low endurance or a need to stop. A patient can therefore be inactive while distressed, inactive while indifferent, or motivated but unable to sustain effort. Those are different clinical problems.

PRESENTATION	CORE BEDSIDE QUESTION	WHAT THE OBSERVER NOTICES	MAIN DISCRIMINATOR
Mania	Has activation increased across several domains?	Accelerated speech, reduced sleep, expansive or irritable behaviour	Excess drive rather than simple agitation
Anxiety	What danger is anticipated, and what is avoided?	Reassurance seeking, vigilance, autonomic distress, reluctance to engage	Fear-linked cognition and behaviour
Apathy	Does action begin without external prompting?	Low initiation, sparse spontaneous activity, reduced interest	Reduced goal direction, not necessarily sadness
Fatigue	Is the wish to act present but effort cannot be sustained?	Slowing, rest seeking, limited endurance	Capacity failure despite intention

MNEMONIC

High drive, high alarm, low drive, low energy

Use that sequence to organise mania, anxiety, apathy and fatigue. It is a starting map, not a substitute for assessing mixed states, cognition, sleep, pain, medication effects and the neurological deficit.

RULE **Describe activation, threat, initiation and endurance separately.** "Not coping" is not a syndrome, and "not participating" does not tell you whether the barrier is fear, absent drive, exhaustion, impaired comprehension or motor disability.

3.2 Post-stroke mania: recognition and diagnosis

Secondary mania may resemble primary mania closely. In a comparative series, elation, pressured speech, flight of ideas, grandiose ideas, insomnia, hallucinations and paranoid delusions occurred with similar frequency in brain-injury-related mania and primary mania when the syndromes were compared. Neurological causation therefore cannot be inferred from an unusual symptom profile. The causal formulation rests instead on the relation to stroke, lesion evidence, age and history, absence or presence of previous mood episodes, exposure to activating substances or medicines, and exclusion of an acute confusional state.

The typical profile reported in the systematic literature was a man with vascular risk, no personal or family psychiatric history and a right-sided cerebral infarct without subcortical atrophy in the reviewed cases. This is a probabilistic profile, not a criterion set. Women, people with psychiatric histories and patients with other lesion patterns still require assessment when the syndrome is convincing. Collateral history is essential because reduced need for sleep, new financial or sexual risk-taking, intrusive sociability and altered premorbid judgement may be clearer to family than to the patient.

Document whether the patient has a full **manic-like episode** or only manic features, and state why stroke is considered causal. The source formulation reserves "with manic features" for presentations that do not meet the full manic syndrome while retaining a stroke-related bipolar formulation. In practice, the note should separate observed behaviour from causal inference: first describe the syndrome, then the temporal and anatomical evidence, then the alternatives considered. This prevents a lesion scan from replacing psychopathology.

TRAP **Writing "right hemisphere equals mania."** Right-sided lesions are an association, not a diagnostic test. The syndrome must still be established clinically, and delirium, disinhibition without mood change, medication effects and primary bipolar illness must be considered.

3.3 Mania: lesion model, course and evidence limits

The exam-salient anatomical contrast is right-sided mania versus the left frontal or basal-ganglia tendency described in post-stroke major depression. In the lesion study, **12 of 17** patients with mania had unilateral right-hemispheric lesions; associated sites included basotemporal and orbitofrontal cortex, frontal white matter, basal ganglia and thalamus across cortical and subcortical lesions. The broad distribution is best understood as a network effect rather than a single "mania centre."

The proposed network is the **right basotemporal-orbitofrontal circuit**, linked to limbic regulation. Direct right basotemporal injury and remote reduction of activity in that region after a subcortical lesion have both been associated with secondary mania. Metabolic imaging found focal right basotemporal hypometabolism in manic patients with right subcortical lesions, supporting a diaschisis account. The conceptual move is important: a structurally distant lesion can disturb mood regulation by disconnecting or suppressing a connected node.

The treatment literature is weak. No randomised controlled treatment study was identified, so apparent efficacy cannot be separated confidently from selection, spontaneous recovery or publication bias in the available evidence. The longitudinal series indicated that most episodes remitted within a few months after post-stroke mania. This should support periodic review of continuing need, not therapeutic passivity during dangerous activation.

PITFALL

Do not treat a new manic syndrome solely from the psychiatric interview. Recheck the neurological course, intercurrent illness, cognition, prescribed and non-prescribed substances, sleep disruption, capacity, vulnerability and immediate behavioural risk. A correct syndrome label can coexist with a wrong causal assumption.

3.4 Post-stroke anxiety: burden, pattern and comorbidity

Anxiety is common enough to require active enquiry, particularly early after stroke. The pooled estimate was highest immediately after the event: 36.7% at 0 to 2 weeks, falling to 24.1% from 2 weeks to 3 months and 23.8% from 3 to 12 months. This pattern supports early screening, but a different meta-analysis suggested a modest rise across later assessment periods. The estimates are

not interchangeable because case definition, setting and measurement method alter who is counted.

That methodological point is visible in the second synthesis: the pooled prevalence was **18% by clinical interview versus 25% by rating scale**. Phobic disorders and generalised anxiety disorder were the most common diagnostic forms in that analysis. Rating scales identify symptom burden; interviews establish whether symptoms form a disorder and whether they are better explained by depression, cognition, communication difficulty or the medical situation. A high scale score is therefore a prompt for formulation, not a self-executing diagnosis.

Anxiety and depression frequently coexist, but neither should be assumed from the other. A review found both **generalised anxiety without depression** and generalised anxiety comorbid with major depression. Setting mattered: community samples and hospital or outpatient samples showed different distributions of anxiety alone and anxiety with depression across the reviewed data. At the bedside, ask separately about fear, worry, avoidance, anhedonia, hopelessness and biological change. Post-stroke depression, including its lesion and treatment evidence, is covered under “Post-stroke depression: epidemiology, lesion correlates, diagnosis and treatment.”

An imaging cohort associated anxiety with right-frontal infarcts even after relevant variables were controlled. This is an association, not a locator. Anxiety can also be sustained psychologically by uncertainty, fear of recurrence, altered bodily sensations and loss of independence. A useful formulation can hold both propositions at once: stroke may disturb circuits involved in emotional regulation, while the meaning and consequences of stroke shape the content of fear.

3.5 Apathy: construct, epidemiology and longitudinal course

Apathy is diminished goal-directed activity arising from reduced motivation. It may be expressed in behaviour, cognition and emotional responsiveness, but the practical bedside sign is lack of self-generated action. The patient may answer correctly, agree that rehabilitation matters and still fail to initiate. Prompting may transiently improve performance because capacity is present but spontaneous mobilisation is weak. This is why an interview conducted entirely through direct questions can miss the syndrome.

Apathy is related to depression but not reducible to it. In the meta-analysis, depression was present in 40.1% of patients with apathy, and the estimated relative risk of depression among apathetic patients was 1.8. The majority of apathetic patients in those data therefore did not have depression. Look for sadness, guilt, hopelessness, negative self-appraisal and distress rather than inferring depression from inactivity alone. Conversely, the absence of visible distress does not make apathy benign.

Course data reinforce the need for repeated assessment. In a longitudinal study, **41.1%** met diagnostic criteria for apathy during the first year after stroke. Mean onset was 3.8 plus or minus 3.3 months, and mean duration was 5.6 plus or minus 2.3 months. Except for a few cases, onset occurred within the first 5 months and the syndrome lasted almost 6 months. A normal early

interview therefore does not close the question, and apparent later "non-compliance" may represent an emerging motivational syndrome.

Older age, depression and cognitive impairment were correlates of apathy, whereas frequency did not differ by sex, stroke type or left versus right hemisphere in the meta-analysis. This non-lateralisation is a useful contrast with the right-sided association described for mania and the right-frontal association described for anxiety. It also warns against using lesion side to dismiss clinically evident apathy.

3.6 Apathy: circuit model and functional importance

The leading anatomical account is disruption of the **anterior cingulate subcortical circuit**. Associated lesions have involved the cingulate region and frontal cortex, ventral striatum, ventral pallidum and magnocellular dorsomedial thalamus. Other reported sites include pons, basal ganglia, dorsal thalamus, posterior limb of the internal capsule and temporal cortex within a distributed circuit model. The shared function is not simple movement but the translation of salience and intention into sustained goal-directed behaviour.

This model explains why apathy can coexist with preserved elementary motor power and why external structure can reveal more ability than spontaneous ward behaviour suggests. It also prevents a moralising interpretation. The person is not necessarily refusing rehabilitation; the neural machinery that assigns value, initiates action and sustains effort may be disrupted. Assessment should therefore compare self-initiated activity with performance after cueing and should obtain observations across settings, not only during a highly structured consultation.

Apathy has prognostic weight. In one cohort, it was a stronger predictor of poor **functional recovery** than depression. A prospective cohort similarly found less recovery in cognition and activities of daily living among apathetic patients during the first year, even after major alternative influences were controlled in the analysis. The clinical implication is direct: recognising apathy changes rehabilitation planning. Goals need to be concrete, initiation may need external scaffolding, progress should be observed rather than inferred from verbal agreement, and caregivers need an explanation that avoids blame.

> **GOOD TO REMEMBER**
>
> Ask the rehabilitation team one question: "What does the patient begin without prompting?" The answer often discriminates apathy more effectively than a polished clinic interview. Then ask what happens after cueing, because improved engagement with structure suggests a target for rehabilitation design.

3.7 Post-stroke fatigue: a disciplined clinical formulation

Post-stroke fatigue should be treated as a real symptom requiring formulation, while resisting unsupported numerical claims. Operationally, ask whether the patient wants to act but experiences effort as excessive, endurance as reduced, or recovery after activity as inadequate. Clarify whether the complaint is mental, physical or both; constant or activity-linked; present on

waking or emerging with exertion; and whether rest restores function. This description is more useful than the word “tired.”

The key distinction from apathy is intention. In fatigue, the patient commonly describes a wish to continue but an inability to sustain output. In apathy, initiation itself is reduced and the lack of activity may not generate equivalent concern. Depression may contribute low energy, but it also requires its own affective and cognitive evidence. Anxiety may consume effort through vigilance and avoidance. Neurological disability may make an ordinary task metabolically and attentively demanding. These formulations can coexist, so forcing the complaint into a single box is often clinically artificial.

Assessment should review the stroke course and functional demands, sleep quality and schedule, pain, nutrition and hydration, intercurrent medical symptoms, medication timing, mood and anxiety, cognition, deconditioning and the balance between activity and rest. Examination and investigations should be directed by the history and neurological context rather than ordered as an undifferentiated panel. Ask the family and therapists whether performance declines predictably, whether the patient can be activated by cueing, and whether apparent fatigue actually follows fear, frustration or cognitive overload.

Management begins by validating the symptom without promising a single cause. Correct remediable contributors, simplify competing demands, pace activity, place important tasks in the patient’s better periods, build rest deliberately into rehabilitation and increase activity according to observed tolerance. Review the response using concrete functional goals. There is no supported prevalence, time-course figure or dedicated treatment trial in the present evidence base, so the academically correct answer is to state the gap rather than manufacture precision.

3.8 Integrated management: locus, focus and modus

Current stroke and liaison-psychiatry guidance should frame care, but the syndrome-specific trial evidence here is limited. Management must remain collaborative with neurology, rehabilitation, general medicine, nursing and caregivers. It should also be phase-sensitive: contain immediate danger first, then restore sleep and participation, then review persistence, recurrence and treatment burden.

Locus. Emergency or inpatient care is appropriate when mania produces dangerous activation, severe behavioural disturbance, inability to cooperate with essential stroke care, marked vulnerability or doubtful decision-making capacity. Medically stable anxiety, apathy and fatigue can often be managed in outpatient or rehabilitation settings, provided risk, function and caregiver capacity are adequate. The locus should change when neurological status, medical stability or behavioural risk changes.

Focus. In the acute phase, target safety, sleep-wake disruption, distress, behavioural containment and participation in essential treatment. Over the short and middle term, target syndrome-specific barriers to rehabilitation: escalating activation in mania, avoidance in anxiety, absent initiation in apathy and limited endurance in fatigue. Longer-term review asks whether the

syndrome has remitted, whether treatment is still needed, whether recurrence risk has changed and whether caregiver burden or social vulnerability persists.

Modus. Combine medical correction, environmental organisation, psychological support, rehabilitation methods, caregiver work and medication when indicated. Use calm structure and reduced stimulation for mania; graded re-engagement and formulation-based psychological work for anxiety; external cueing, scheduled activity and meaningful goals for apathy; and pacing with correction of contributors for fatigue. Preserve autonomy while matching supervision to capacity and risk. For antidepressant pharmacology and drug selection, see the Mood disorders: pharmacotherapy notes.

For mania, case data support a trial of lithium in selected patients, but non-response to lithium and carbamazepine has also been reported, and no randomised treatment evidence exists for post-stroke mania. Choice and monitoring therefore require individual assessment of medical status, interactions, cognition, adherence and rehabilitation goals. Do not copy a primary-bipolar algorithm into a medically complex stroke survivor without reconsidering risk and the likely course.

For anxiety, merged double-blind data comparing nortriptyline with placebo in comorbid anxiety and depression found anxiety ratings improved more rapidly than depressive ratings with nortriptyline. However, the systematic review conclusion was that evidence remained insufficient to guide treatment of post-stroke anxiety. Use that evidence to justify humility, not nihilism: treat clinically important anxiety, define the target, choose an intervention compatible with the patient's medical and cognitive context, and review actual benefit.

For apathy, nortriptyline, bromocriptine, methylphenidate, amantadine, selegiline and tacrine have been used with some success. Nefiracetam at **900 mg/day for post-stroke apathy** outperformed placebo on apathy ratings in a double-blind trial, but it is not approved by the FDA for any indication. These data do not create a routine prescription hierarchy. Off-label treatment requires explicit targets, risk review, informed discussion and withdrawal when meaningful functional benefit is absent.

- Define the target syndrome and the behaviour that will show improvement.
- Record neurological, medical, cognitive, psychological and social contributors.
- Match supervision and treatment setting to risk, capacity and rehabilitation needs.
- Review function with the patient, caregiver and rehabilitation team, not symptoms in isolation.

After stroke, reduced activity is not synonymous with depression: separate fear, loss of initiation and loss of endurance before treating the label.

4 · Post-stroke psychosis, emotional incontinence and denial of deficit

Why these syndromes belong together. After stroke, disturbed reality testing, involuntary emotional display, collapse under challenge and unawareness of disability may all be reported as "behavioural change." That phrase is clinically inadequate. Each phenomenon concerns a different psychological operation: psychosis concerns belief or perception; emotional incontinence concerns control of expression; the catastrophic reaction concerns coping when impaired performance is exposed; and anosognosia concerns awareness of deficit. The marks lie in defining the operation that has failed, then relating it to time course, lesion pattern and function without turning an association into a diagnostic rule.

IN INDIA

Reserve antipsychotic treatment for genuine psychotic risk and use the lowest effective dose, since older stroke patients tolerate these agents poorly and are prone to falls. No specific agent is established for post-stroke psychosis, so choose cautiously and review early.

These distinctions also determine bedside conduct. A fixed false belief requires assessment of its content and consequences. Incongruent laughing or crying requires separation of outward display from inner mood. An outburst during a difficult task may reflect overload rather than sustained hostility. A patient who denies paralysis cannot be assumed to understand safety advice merely because language is fluent. Collateral observation, longitudinal review and direct comparison between what the patient says and does are therefore central throughout this section.

0.4% to 3.1%	6.1 MONTHS	19%	32.3%
REGISTRY FREQUENCY OF POST-STROKE PSYCHOSIS	MEDIAN ONSET AFTER STROKE IN REGISTRY DATA	CATASTROPHIC REACTION IN AN ACUTE-STROKE SURVEY	ANOSOGNOSIA WITH COMPLETE CONTRALATERAL MOTOR IMPAIRMENT

4.1 The syndrome map

The safest opening question is not "Is this psychiatric?" but "What exactly is abnormal?" A syndrome may be apparent in the content of thought, the form of emotional expression, the response to task failure or the appraisal of bodily function. This descriptive step prevents several common category errors. Crying does not automatically establish depressed mood; shouting during testing does not automatically establish a persistent aggressive syndrome; denial of paralysis is not simply poor motivation; and a delusion is not explained merely by the presence of a brain lesion.

SYNDROME	DISTURBED OPERATION	CHARACTERISTIC BEDSIDE CLUE	KEY DISCRIMINATION
Post-stroke psychosis	Reality appraisal	Delusion or hallucination emerging in relation to stroke	Establish psychosis, then assess neurological attribution and competing causes
Emotional incontinence	Control of emotional expression	Involuntary laughing or crying that does not match reported inner feeling	Expression is not equivalent to mood
Catastrophic reaction	Coping under challenge	Brief outburst when a deficit is exposed by a demanding situation	Provoked overload is not the same as sustained aggression
Anosognosia	Awareness of deficit	Statements and behaviour show failure to recognise impairment	Unawareness differs from indifference and from spatial inattention
Anosodiaphoria	Emotional concern about deficit	The deficit is admitted but treated with striking unconcern	Knowledge is present; appropriate concern is reduced

MNEMONIC

Reality, expression, overload, awareness

Use this sequence to localise the clinical problem: reality is altered in psychosis; expression escapes voluntary control in emotional incontinence; overload provokes the catastrophic reaction; awareness fails in anosognosia. Anosodiaphoria sits beside awareness because the deficit is known but not regarded with expected concern.

RULE **Name the disturbed operation before naming the diagnosis.** The same visible behaviour can arise from altered belief, uncontrolled expression, task-related distress or absent insight, and management follows the mechanism suggested by that distinction.

4.2 Post-stroke psychosis: phenotype and epidemiology

Post-stroke psychosis is rare and is identified principally by delusions or hallucinations arising after stroke. Registry estimates vary substantially: hospital studies reported frequencies of 0.4% and 3.1%, with an incidence of **1.1 per 1,000 person-years**. The reported median interval from stroke to onset was 6.1 months. Delayed emergence is therefore entirely compatible with the syndrome; a symptom-free acute admission does not close the later diagnostic window.

The broader systematic review included **76 studies** and found an average age of 66.6 years. Delayed onset was common. **Delusional disorder** was the most frequent psychotic presentation, followed by schizophrenia-like psychosis and mood disorder with psychotic features. Estimated prevalences were 4.67% for delusions and 5.05% for hallucinations; the reported 12-year incidence was 6.7%. These figures come from different study designs and constructs, so they should be presented as complementary estimates rather than forced into a single rate.

Phenomenology still comes before attribution. Describe the belief or perceptual experience in ordinary psychiatric terms: its content, conviction, preoccupation, associated affect, behavioural consequences and effect on care. Ask whether the patient acts on it, whether it concerns staff or family, whether it impairs eating, medication acceptance or rehabilitation, and whether collateral observers confirm a change from baseline. A neurological event does not make psychopathology self-explanatory.

A practical assessment sequence is to establish whether a delusion, hallucination or psychotic mood syndrome is actually present; reconstruct onset in relation to the stroke, recovery setting and later transitions of care; seek collateral evidence about premorbid beliefs, previous episodes, behaviour and functional consequences; and assess immediate risk arising from fear, command experiences, misidentification, refusal or unsafe action.

4.3 Lesion model, causal formulation and prognosis

The recurring anatomical association is the right hemisphere. Reported lesions have involved right frontal, temporal and parietal regions and the right caudate nucleus. In a study of five patients, every patient had a right-hemispheric lesion, mainly in frontoparietal regions. Neuroimaging analyses further associate delusional belief with the **right lateral prefrontal cortex**, either through a lesion in that region or disruption of connectivity to it.

This evidence supports a network formulation rather than a crude "psychosis centre." A focal lesion may disturb appraisal, belief evaluation and the integration of internally generated experience with external evidence. Connectivity findings matter because the decisive functional disruption need not be confined to the visibly damaged tissue. Right-hemispheric injury, particularly involving or disconnecting lateral prefrontal systems, together with seizures and subcortical brain atrophy, has been implicated in pathogenesis. The association increases plausibility; it does not establish causation in an individual.

Causal formulation should therefore integrate syndrome, chronology, neurological evidence and alternatives. Consider an acute confusional state, a mood syndrome with psychotic features, a primary psychotic illness, seizure-related phenomena, substance or medication effects, sensory impairment and a pre-existing cognitive disorder. These are differential labels here, not substitute mini-chapters. The clinician should document which observations support stroke attribution and which uncertainty remains.

TRAP **"Right hemisphere lesion equals post-stroke psychosis."** The laterality and network findings are associations. Diagnosis still requires psychotic phenomenology, a defensible temporal and causal relation, and active exclusion of better explanations.

Prognosis deserves explicit weight. Post-stroke psychosis has been associated with poor functional outcomes and high mortality, adding substantially to the burden of stroke. That association should prompt careful follow-up and coordination, but it should not be converted into deterministic counselling. Functional outcome is shaped by neurological disability, psychiatric symptoms, treatment engagement, environment and support. Record psychiatric and

neurological trajectories separately so improvement in a psychiatric domain is not mistaken for neurological recovery.

4.4 Emotional incontinence in the post-stroke setting

Emotional incontinence denotes involuntary, uncontrollable laughing or crying that may arise spontaneously or after minor provocation and is characteristically unrelated to the patient's reported internal emotional state. Other names include pathological laughter and crying, emotional lability, pathological emotionalism, involuntary emotional expression disorder and pseudobulbar affect. No established DSM-5 diagnostic criteria are provided for the phenomenon.

The core clinical move is to separate display from experience. Observe the episode, then ask how the patient felt before, during and after it. Explore whether the expression was wanted, controllable, congruent with thought content and proportionate to the trigger. Emotional incontinence can be distressing and socially disabling even when the patient denies sustained sadness or mirth. Conversely, genuine depressed mood may coexist with involuntary crying. The visible act alone cannot settle the mood diagnosis.

Post-stroke estimates vary with timing and informant method. Among 508 consecutive patients with ischaemic stroke, pathological emotions affected 9.4% during acute admission and 11.7% three months later. Another study of 516 patients assessed three months after stroke estimated pseudobulbar affect at 7.6%. In 127 patients assessed at the same post-stroke interval, prevalence was **17.9% with caregiver involvement versus 6.3% by patient report alone**.

> **GOOD TO REMEMBER**
>
> Ask the caregiver what the episode looked like, what triggered it and whether it matched the patient's stated feeling. The large informant-related difference in observed prevalence shows why patient report alone can miss involuntary emotional expression after stroke.

This section owns the post-stroke context and recognition of emotional incontinence. Its detailed phenomenology, pathophysiology and treatment are covered under "Pseudobulbar affect: phenomenology, pathophysiology and treatment." Keeping that boundary clear also prevents treatment material from being duplicated or reconstructed from memory.

4.5 Catastrophic reaction: failure under challenge

Catastrophic reaction, first described by Kurt Goldstein, is a provoked emotional outburst containing varying mixtures of anger, frustration, depressed affect, tearfulness, refusal, shouting, swearing, displacement, renunciation, aggression or compensatory boasting. The original account interpreted it as failure of the person to cope when confronted with a serious physical or cognitive defect. Stressful situations such as cognitive testing commonly provoke it, and it usually settles within a few minutes. No formal DSM-5 diagnostic criteria are established.

The time-locked relation to challenge is the bedside clue. The patient may begin a task, encounter evidence of impaired performance and rapidly shift into distress, refusal or anger. The response may look disproportionate if the observer attends only to the immediate test item; it becomes intelligible when understood as exposure of lost capacity. The examiner should stop escalating

demand, acknowledge distress, simplify communication and return later with a graded task. Arguing about the behaviour while the patient is overwhelmed usually supplies more challenge rather than more understanding.

In a survey of 62 consecutive patients with acute stroke using the Catastrophic Reaction Scale, 12 patients, or 19%, showed catastrophic reactions. Of those affected, 9, or 75%, met criteria for major depression and a further 2, or 17%, had minor depression. A larger prospective study assessed 202 consecutive patients within 4 days of stroke; anger occurred in 71 patients, or 35%, and 26 were severely angry. These findings make depressive assessment necessary, but the reaction itself should not be relabelled automatically as depression.

Lesion evidence points towards anterior subcortical systems. Patients with catastrophic reactions had more basal-ganglia lesions than acute-stroke controls. Among depressed patients, those with a catastrophic reaction had more anterior, mainly subcortical lesions: 8 of 9 had subcortical lesions compared with 3 of 9 depressed patients without the reaction. Laterality was not specified in that evidence and should not be invented.

4.6 Anosognosia and anosodiaphoria

Anosognosia, a term introduced by Joseph Babinski, originally referred to lack of awareness of hemiplegia. Its use has extended to unawareness of other post-stroke deficits, including cortical blindness, hemianopia and amnesia. The deficit is behavioural as well as verbal: a patient may deny weakness, attempt an impossible action or plan discharge without accommodating the impairment. A polished explanation of the stroke does not prove awareness if behaviour contradicts it.

Anosodiaphoria is different. Here the patient admits the deficit but appears unconcerned about it. The distinction is between absent knowledge and altered concern. It is also important not to merge anosognosia with hemispatial neglect. Neglect concerns allocation of spatial attention and is covered in the dedicated parietal-lobe syndromes section; anosognosia concerns awareness of the deficit itself. The syndromes may be encountered together, but they are not synonyms.

The evidence base is heterogeneous and timing dependent. A meta-analysis by Lorenzo Pia and colleagues identified 52 studies conducted from 1938 to 2001. Frequency of anosognosia for hemiplegia ranged from 20% to 44%, depending on when assessment occurred relative to stroke. Among patients with complete contralateral motor impairment, prevalence was **32.3%**. Anosognosia is uncommon beyond 3 months after stroke, so repeated reassessment is appropriate rather than treating the acute finding as a fixed personality trait.

Laterality is exam-salient but not absolute. Denial of hemiplegia is associated particularly with right-hemisphere damage producing denial of left-sided paralysis. The underlying functions are not strongly lateralised, yet deficits are more common and more severe after right-hemisphere injury. The correct formulation is therefore “right-sided predominance,” not “right-sided requirement.”

4.7 Bedside assessment and differential reasoning

Triangulate words, behaviour and collateral evidence. Direct questioning alone is especially weak when awareness, expression or coping is the very function under examination. Ask the patient to describe the stroke and its consequences, then observe whether choices and attempted actions incorporate those consequences. Ask relatives and rehabilitation staff about change from baseline, episodes outside the interview and the conditions that trigger them. Repeat assessment because psychosis may emerge after delay, emotional displays may be under-reported, catastrophic reactions are situation dependent and anosognosia often changes with time.

- **For psychosis:** document belief or perception, conviction, distress, action and relation to the neurological course.
- **For emotional incontinence:** compare outward display with reported inner affect, trigger, voluntariness and control.
- **For catastrophic reaction:** reconstruct the task demand, moment of failure, emotional sequence and recovery after the challenge ends.
- **For anosognosia:** compare verbal acknowledgement with practical behaviour, safety decisions and response to demonstration of the deficit.
- **Across syndromes:** assess communication, cognition, sensory impairment, mood, an acute confusional state, medication or substance effects and the possibility of coexisting syndromes.

> **PITFALL**
>
> Do not confront denial by repeatedly forcing proof of failure. Demonstration may be informative, but escalating challenge can provoke distress or a catastrophic reaction and may damage engagement. Use graded observation, collateral information, environmental safeguards and repeated review.

Capacity and awareness must also be kept conceptually separate. A patient may lack awareness of a focal deficit yet reason adequately about an unrelated decision; another may verbally admit weakness while failing to appreciate its practical consequences. Decision-making should therefore be assessed for the actual choice at hand, using communication support and concrete information. The label “poor insight” is too broad to replace this analysis.

4.8 Management framework: locus, focus and modus

Locus. Care should sit where neurological stability, behavioural risk and rehabilitation need can all be managed. Emergency or inpatient assessment is appropriate when there is acute neurological change, severe behavioural disturbance, inability to maintain safety, refusal of essential care or an acute confusional presentation. Medically stable patients may be followed through coordinated stroke, rehabilitation and mental health services. The setting should change when risk, support or neurological status changes.

Focus. In the acute phase, clarify the syndrome, identify urgent mimics, protect the patient and staff, and preserve participation in stroke care. In the short term, reduce distress and dangerous behaviour, explain the syndrome to caregivers, adapt rehabilitation demands and document

functional consequences. Through later recovery, reassess emerging psychosis, persistence of emotional episodes, change in deficit awareness and caregiver burden. Psychiatric, neurological and functional outcomes should be tracked on separate axes.

Modus. Management is multimodal. Use calm communication, predictable routines, correction of reversible contributors, environmental safety, rehabilitation adaptations, caregiver education and psychological support. For catastrophic reactions, grade task difficulty, stop when overload escalates and return after recovery. For anosognosia, build safety around demonstrated behaviour rather than relying on verbal agreement. For emotional incontinence, explain that expression may not represent inner mood and use the dedicated pseudobulbar affect section for mechanism and treatment. For psychosis, assess risk and impairment, coordinate with neurology and consult current stroke and prescribing guidance before patient-specific pharmacological decisions.

No verified agent or dose is available here for post-stroke psychosis, so a specific regimen should not be reconstructed from general psychosis practice. The same restraint applies to formal guideline attribution: where the evidence does not supply a named authority, use current local guidance rather than attaching a remembered label. This is not therapeutic nihilism. It is the correct separation between well-grounded syndrome recognition and patient-specific prescribing that requires a broader evidence base.

The integrated answer is concise in logic even when detailed in execution: identify which psychological operation is disturbed; describe phenomenology before inferring cause; preserve the delayed course and right-hemisphere associations of psychosis; distinguish involuntary expression from mood; recognise a brief challenge-linked catastrophic reaction and its close association with depression; and separate unawareness from unconcern and from neglect. The purpose is not merely accurate naming. It is safer care, better rehabilitation and a formulation that respects what the patient can perceive, control and understand.

> After stroke, distinguish altered reality, involuntary expression, challenge-linked overload and unawareness of deficit; the visible behaviour is not the diagnosis.

5 · Vascular cognitive impairment: spectrum, mechanisms and criteria

The marks lie in separating a vascular spectrum from the old dementia-only stereotype. **Vascular cognitive impairment** is the umbrella for cognitive impairment attributable to pathology of cerebral blood vessels; vascular dementia is its more severe form, not a synonym for the whole spectrum. The vascular injury may be ischaemic or haemorrhagic, although ischaemic disease dominates the literature because it is much more prevalent. A strong answer therefore moves from spectrum to substrate, phenotype and course, then tests whether cognitive impairment, cerebrovascular disease and their association have all been demonstrated.

IN INDIA

The achievable win is prevention: treat hypertension, diabetes and other vascular risk in primary care, and prioritise sustained risk-factor control over any expectation of structured cognitive rehabilitation.

The second source of marks is criteria discipline. Diagnostic frameworks share a causal core but set different evidentiary thresholds. A patient may fit a permissive clinical category yet fail a stricter research definition. The candidate must state each framework on its own terms, use structural imaging as evidence rather than as a label generator, and allow mixed pathology when the clinical and radiological story is not pure. The general dementia syndrome, cognitive instruments and full reversible-cause workup belong in the Dementias and neurocognitive disorders notes.

RULE **Make three links before naming vascular causation.** Establish cognitive impairment, cerebrovascular disease, and a defensible association between them. Different criteria operationalise those links differently, but none permits vascular disease on a scan to stand in for the cognitive syndrome.

5.1 The spectrum: impairment before dementia

The preferred conceptual move is from a categorical endpoint to a continuum. **Vascular cognitive impairment without dementia** names the pre-dementia stage explicitly. It allows recognition of vascular cognitive decline before the functional threshold for the severe syndrome has been crossed. At the other end lies vascular dementia. Between them are clinically important states in which cognition has declined, the vascular contribution is plausible, but the person does not have the severe form.

This terminology changes clinical reasoning. If “vascular dementia” is the only category considered, early executive slowing or attention difficulty may be dismissed until disability becomes conspicuous. The broader term asks a different question: is vascular brain disease already impairing cognition, regardless of whether the endpoint called dementia has been reached? It also makes longitudinal formulation easier because a patient can move along the spectrum as vascular burden, cognitive performance and functional independence change.

CONCEPT	WHAT IT CONTRIBUTES	WHAT IT DOES NOT ESTABLISH BY ITSELF
VCI without dementia	Recognises the pre-dementia end of the vascular cognitive spectrum.	It does not erase the need to associate cognition with vascular disease.
Vascular dementia	Names the more severe end of vascular cognitive impairment.	It should not replace VCI as the umbrella term.
Mixed vascular and Alzheimer pathology	Allows concurrent vascular and degenerative contributions.	A vascular lesion need not be the sole cause of decline.
Vascular depressive disorder	Places late-life affective and executive change beside the cognitive spectrum.	It is not automatically vascular neurocognitive disorder.

A spectrum is not an invitation to diagnose by impression. The breadth lies in severity and presentation, while causal attribution remains disciplined. The clinician still needs evidence that cognition has changed, that cerebrovascular pathology is present, and that they belong in the same formulation. General diagnostic thresholds for major or mild neurocognitive disorder should be taken from the Dementias and neurocognitive disorders notes; this section supplies the vascular lens.

5.2 Substrates, mechanisms and the shape of decline

Vascular cognitive injury is heterogeneous because the vascular lesion may be large-vessel or microvascular, and focal, multifocal or diffuse. These patterns can coexist. The anatomical substrate therefore determines which network fails, while lesion accumulation and recurrence influence the course. Strategic infarction, multi-infarct burden and small-vessel disease are named in the accompanying figure; their detailed anatomy and inherited small-vessel syndromes are covered in the dedicated strategic-infarct, small-vessel disease, post-stroke dementia and CADASIL section.

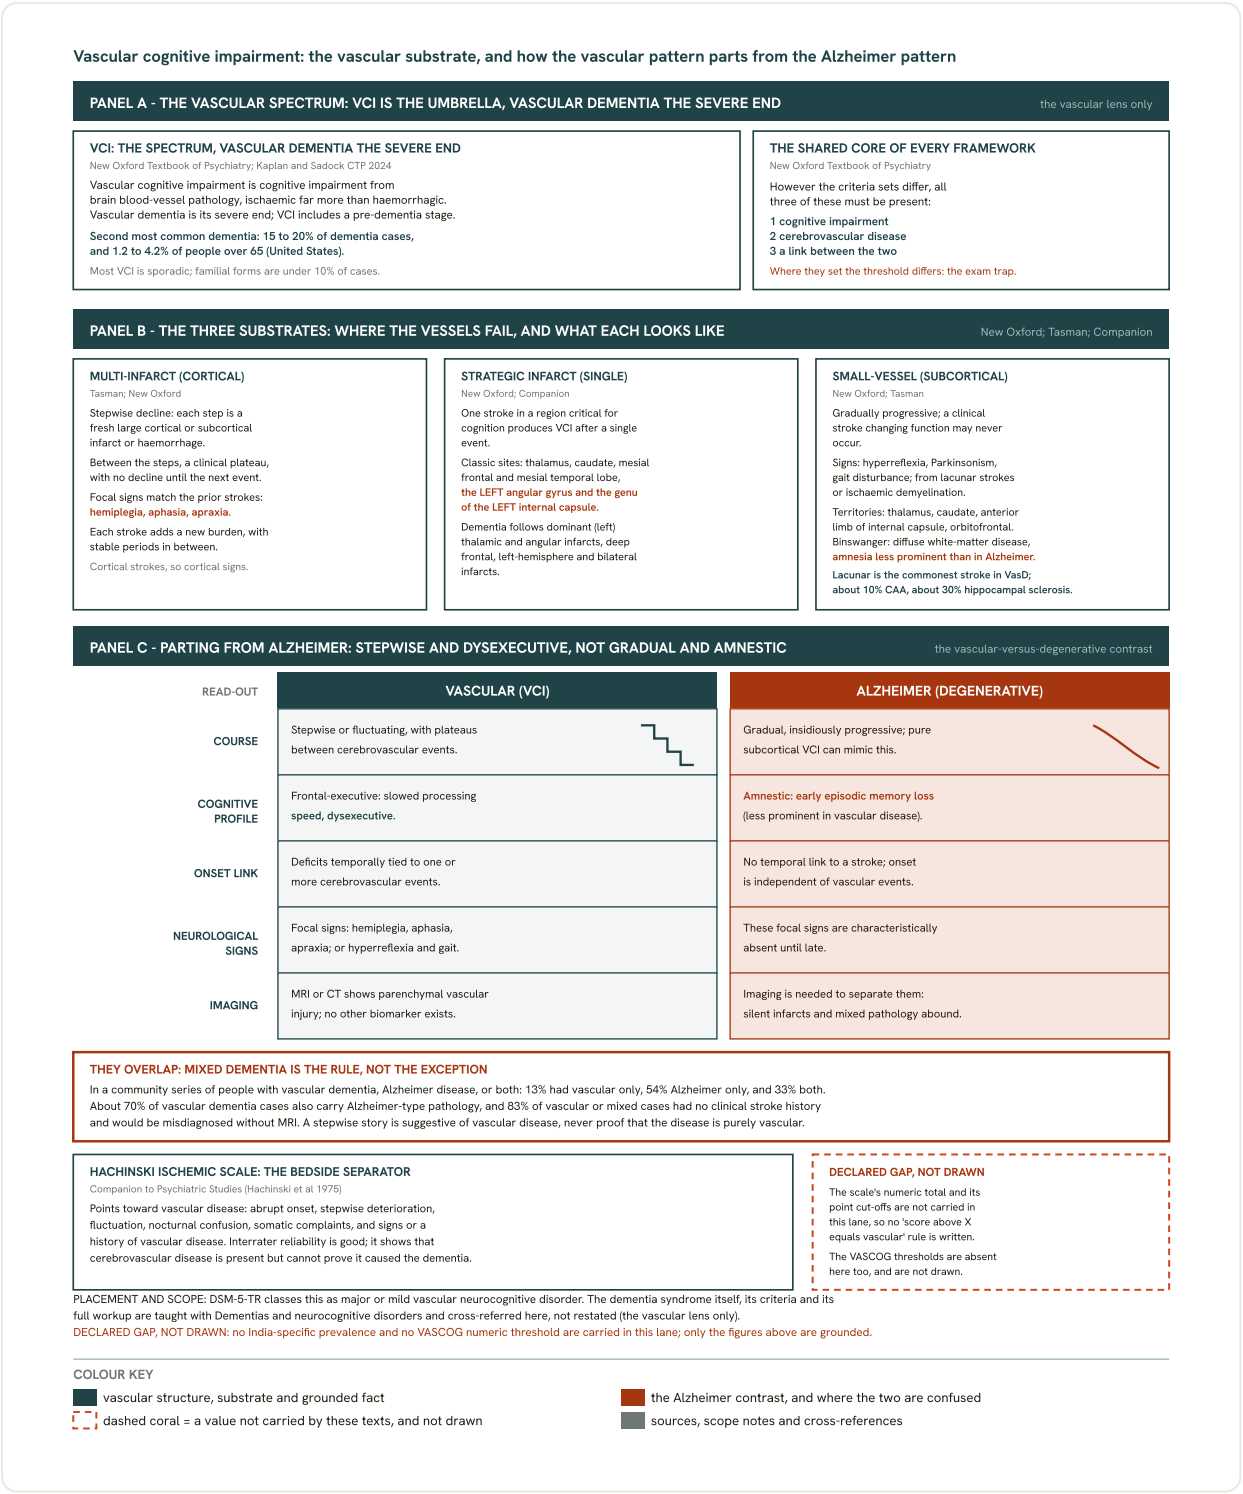


Fig 34.2 **Vascular cognitive impairment, and how it parts from Alzheimer disease.** VCI is the umbrella and vascular dementia its severe end, sitting on three substrates: multi-infarct disease deteriorates stepwise with plateaus and focal signs, a strategic single infarct in a critical site such as the left angular gyrus or the genu of the left internal capsule produces impairment after one event, and small-vessel disease creeps forward with hyperreflexia and gait change and often no clinical stroke. Against Alzheimer disease the vascular pattern reads stepwise and dysexecutive rather than gradual and amnestic, yet mixed pathology is the rule and imaging is what separates them; the Hachinski scale weighs the bedside features while its numeric cut-off is a declared gap. (New Oxford Textbook of Psychiatry; Kaplan and Sadock CTP 2024; Tasman Psychiatry; Companion to Psychiatric Studies; DSM-5-TR.)

The common pathophysiological endpoint is injury to brain parenchyma, reached through several interacting routes. **Hypoperfusion and hypoxia** reduce the delivery needed for network function and tissue survival. Oxidative stress and inflammation promote endothelial dysfunction.

Injury to the vascular control system can then impair autoregulation and neurovascular coupling, weakening the matching of blood flow to neural demand. The point is not to memorise isolated arrows but to see a feed-forward vulnerability: damaged vessels regulate perfusion poorly, and poorly supported neural networks become less resilient.

Course follows burden and distribution rather than one compulsory graph. Many patients with multiple infarctions show **stepwise or fluctuating decline**, with intervening stability and sometimes improvement. Other recognised courses include gradual onset with slow progression, rapid development followed by relative stability, and more complex patterns. Pure subcortical vascular neurocognitive disorder may progress slowly enough to simulate neurocognitive disorder due to Alzheimer disease.

MNEMONIC

V-A-S-C-U-L-A-R

Vessels establish the substrate; **A**ttention and executive function may dominate; **S**teps are classic but not compulsory; **C**ognition must be demonstrated; **U**neven profiles may occur; **L**ink the lesion to the decline; **A**llow mixed pathology; **R**equire a better explanation to be excluded.

5.3 Cognitive phenotype: executive is classic, not universal

The cognitive profile varies with the location and combination of vascular lesions. The classic pattern is **frontal-executive dysfunction**; subcortical ischaemic white-matter disease may make this especially prominent. Reduced processing efficiency, difficulty organising a sequence, poor mental flexibility and impaired goal maintenance can therefore be more revealing than a memory-led interview. This is also why a collateral description such as “he remembers but cannot manage anything” deserves domain-based analysis rather than being dismissed as inconsistency.

Classic does not mean obligatory. Other cognitive patterns occur because lesion location determines the deficit. A focal lesion can produce a sharply patterned loss, multifocal disease can yield an uneven profile, and diffuse disease can create broad inefficiency. The clinician should describe the domains actually affected and resist converting “vascular” into a fixed neuropsychological template.

At assessment, establish change from the person’s previous cognitive and functional baseline, and look specifically for complex attention, processing speed and executive organisation rather than asking only about memory. Reconcile the domain pattern with temporal course, neurological findings and the distribution of vascular injury. Finally, ask whether another brain disease or systemic disorder explains the presentation better.

The vascular-versus-degenerative contrast is therefore probabilistic, not absolute. Stepwise decline with focal neurological change strongly supports a vascular account, but gradual progression does not exclude it. Conversely, an executive profile is compatible with vascular disease but cannot prove its cause. The accompanying figure should be read as a contrast of typical patterns and substrates, not as a binary bedside algorithm.

PITFALL

Do not diagnose vascular cognitive impairment because a report mentions white-matter change. Imaging establishes vascular burden; the diagnosis still requires demonstrated cognitive impairment and a clinically credible association between cognition and vascular disease.

5.4 Epidemiology and the mixed-pathology reality

In United States data, vascular dementia is the second leading cause of dementia after Alzheimer disease. It affects **1.2-4.2% of people older than 65 years** and accounts for **15-20% of dementia cases**. Prevalence rises with age, sex differences remain unclear, and reliable prevalence statistics for VCI below the dementia threshold are lacking. These figures must retain their geographical scope; they are not Indian prevalence estimates.

13%	54%	33%	83%
VASCULAR ONLY	ALZHEIMER ONLY	BOTH PATHOLOGIES	NO STROKE HISTORY IN VASCULAR OR MIXED GROUP

Mixed pathology is not a marginal footnote. In a large community study of people with vascular dementia, Alzheimer disease or their combination, 13% had vascular disease alone, 54% Alzheimer disease alone and 33% both. Among those with vascular or mixed vascular-Alzheimer disease, **83%** had no clinical history of stroke and might have been misdiagnosed without MRI. Separately, some studies report Alzheimer-type pathology in 70% of people with vascular dementia. The exact estimates come from different studies and should not be collapsed into one denominator.

Secular data add a prevention perspective. In the Framingham Heart Study, overall dementia incidence declined over 40 years, falling by 20% every 10 years, while vascular dementia declined faster than Alzheimer disease. The excess risk of dementia after stroke fell from nine times that of people without stroke to **less than twice**. This trend supports the modifiability of population burden without implying that established vascular injury is uniformly reversible.

United States disparity data also warn against treating vascular risk as purely biological. Progression from ischaemic stroke to vascular neurocognitive disorder within 5 years was almost twice as likely among African Americans as among non-Latinx White people and occurred at younger ages. Higher rates of hypertension and diabetes, limited formal education and low socioeconomic status were proposed contributors. The finding is context-specific but illustrates how vascular exposure and social determinants can converge.

5.5 The diagnostic core and the role of imaging

Across frameworks, the shared diagnostic core is stable: **cognitive impairment**, evidence of cerebrovascular disease, and evidence that the vascular disease is associated with the cognitive syndrome. Each pillar answers a different question. Cognitive evidence establishes the clinical target. Vascular evidence establishes a plausible substrate. Association turns coexistence into an aetiological formulation.

Association should be argued, not declared. Useful reasoning includes temporal correspondence with a cerebrovascular event, concordance between lesion distribution and the affected cognitive domains, a course compatible with the observed vascular burden, and the absence of a better explanation. No single feature is universally mandatory across all systems. The weight of each feature depends on the framework being applied.

Structural neuroimaging with MRI or CT has an important diagnostic role. It can demonstrate parenchymal injury, reveal clinically silent vascular burden and help judge whether the distribution is sufficient to account for the deficits. Yet a lesion remains evidence of cerebrovascular disease, not proof of causation.

There are no other established biomarkers for major or mild vascular neurocognitive disorder. Diagnosis therefore remains an integration of history, examination, cognitive assessment and imaging rather than a laboratory confirmation. The general blood investigations, cognitive instruments and reversible-cause assessment are not repeated here; they are covered in the Dementias and neurocognitive disorders notes and the Psychotherapies, clinical assessment, MSE and nosology notes.

GOOD TO REMEMBER

A normal history for overt stroke is not a normal vascular history. Ask about abrupt functional changes and focal symptoms, examine for neurological signs, and inspect imaging before concluding that gradual decline is purely degenerative.

5.6 DSM-5-TR: major or mild vascular neurocognitive disorder

DSM-5-TR major or mild vascular neurocognitive disorder begins by requiring that the criteria for major or mild neurocognitive disorder are met. Vascular aetiology is then suggested through an either-or route: cognitive onset is temporally related to one or more cerebrovascular events, or decline is prominent in complex attention, including processing speed, and frontal-executive function. This alternative phenotype route matters because a documented stroke is not compulsory at this point in the criteria.

The remaining requirements restore causal discipline. There must be cerebrovascular disease evident from history, physical examination and/or neuroimaging that is sufficient to account for the deficits. The presentation must also not be better explained by another brain disease or systemic disorder. Thus, executive decline alone cannot establish the diagnosis, and incidental vascular lesions cannot carry the entire inference.

The probability specifier asks how firmly the vascular attribution is supported:

- Probable vascular neurocognitive disorder may be assigned when clinical criteria are supported by neuroimaging evidence of significant vascular parenchymal injury.
- It may also be probable when the syndrome is temporally related to one or more documented cerebrovascular events.

- A further probable route requires both clinical and genetic evidence of cerebrovascular disease, with CADASIL given as an example; the inherited disorder itself is covered in the dedicated substrate section.
- Possible vascular neurocognitive disorder applies when clinical criteria are met but imaging is unavailable and a temporal relationship to cerebrovascular events is not established.

The logic is deliberately more permissive than criteria designed to create homogeneous research samples. It lets a clinician preserve a vascular hypothesis when the syndrome and phenotype fit but decisive imaging or event linkage is missing. "Possible" records uncertainty; it does not authorise omission of the clinical criteria or the exclusion of better explanations.

5.7 Competing frameworks and the examination trap

FRAMEWORK	THRESHOLD OR EMPHASIS	PRACTICAL CONSEQUENCE
DSM-5-TR	Temporal relation to cerebrovascular events or prominent complex-attention and frontal-executive decline; vascular disease sufficient to explain deficits.	Possible vascular NCD remains available without imaging and without an established temporal link.
NINDS-AIREN research criteria	Clinical and neuroimaging evidence of cerebrovascular disease, plus a temporal relation to dementia.	Stricter case selection; high specificity but low sensitivity was found when clinical criteria sets were compared.
DSM-IV(R)	Dementia with focal neurological signs and symptoms, or aetiologically relevant cerebrovascular disease.	Historical either-or route.
ICD-10	Focal brain damage, aetiologically relevant cerebrovascular disease and an uneven distribution of cognitive deficits.	Adds the patchy cognitive-profile requirement.
Hachinski Ischemic Scale	Abrupt onset, stepwise deterioration, fluctuation, nocturnal confusion, somatic complaints and vascular risks or signs.	May support the presence of significant cerebrovascular disease but cannot demonstrate causation.

TRAP **Do not fuse permissive clinical criteria with strict research criteria.** DSM-5-TR permits possible vascular NCD when imaging is unavailable and no temporal link is established. NINDS-AIREN requires both clinical and imaging evidence plus a temporal relationship. A gradual executive presentation that otherwise meets the clinical criteria can therefore enter the former category yet fail the latter threshold.

The NINDS-AIREN criteria are widely used in research. Their restrictive design reduces false-positive inclusion, but the reported combination of high specificity and low sensitivity means genuine clinical cases may be missed. That trade-off is appropriate for some research questions and too restrictive for others. It is not evidence that one framework has discovered the single correct boundary.

The Hachinski Ischemic Scale is historically important and can show good interrater reliability, but it has not been well validated. Its listed clinical features suggest vascular burden; they do not prove that vascular disease caused the cognitive syndrome. Numeric decision thresholds are intentionally not reproduced because the available source does not establish them.

5.8 Clinical synthesis, differential boundaries and answer structure

A clinically useful formulation proceeds in a fixed sequence. Define the level on the vascular cognitive spectrum. Describe the affected cognitive domains and functional change. Characterise onset and course. Identify neurological and imaging evidence of cerebrovascular disease. Explain the temporal or anatomical association. Then assess competing and coexisting causes, especially mixed vascular and degenerative pathology.

The affective boundary deserves attention. The **vascular depression hypothesis** proposes that small-vessel cerebrovascular disease can damage corticostriatal circuits and create a late-onset depressive subtype. In older adults with progressive small-vessel ischaemic disease, depressive symptoms may occur with psychomotor retardation and executive dysfunction. This "vascular depressive disorder" must be differentiated from vascular neurocognitive disorder, and widely accepted diagnostic criteria remain lacking. Mood symptoms and executive inefficiency can therefore be adjacent manifestations without being interchangeable diagnoses.

The examination answer should close by naming uncertainty precisely:

- "Possible" is an evidentiary category within DSM-5-TR, not a substitute for the clinical syndrome; mixed pathology may be more credible than forcing vascular and Alzheimer disease into mutually exclusive boxes.
- Absent overt stroke does not exclude vascular injury, and gradual progression or memory complaints do not exclude VCI; concordance among phenotype, course and vascular burden remains decisive.

For a long answer, begin with the umbrella definition and pre-dementia spectrum, explain vascular mechanisms and variable course, state the classic but non-universal executive profile, and then compare criteria on shared axes. Give DSM-5-TR as the contemporary clinical cross-map, contrast it explicitly with NINDS-AIREN and historical ICD-10, and end with imaging, mixed pathology and causal restraint. This preserves bedside usefulness as well as research-framework accuracy.

> Vascular cognitive impairment is a spectrum, not a synonym for vascular dementia: prove cognition, prove cerebrovascular disease, link them, and state which criteria threshold you are applying.

6 · Strategic infarct, small-vessel disease, post-stroke dementia and CADASIL

The marks lie in pattern recognition. A candidate must distinguish cognitive decline caused by accumulated large infarcts, a single lesion in a cognitively critical site, diffuse small-vessel injury, and an inherited arteriopathy. The decisive clues are the relation to stroke, the tempo between

events, the neurological examination, the distribution of lesions and the family history. The same final label can conceal very different vascular substrates.

IN INDIA

Base the working assessment of suspected inherited small-vessel disease on the clinical picture and imaging, and refer for confirmatory genetic testing only when the result would change management.

This section develops those substrates rather than restating the general spectrum, mechanisms or diagnostic criteria of vascular cognitive impairment. It also separates post-stroke dementia as an outcome from any one pathological subtype: a patient may decline after repeated infarcts, after one strategically placed infarct, or through an apparently silent subcortical process. The full dementia syndrome and its general workup are covered in the Dementias and neurocognitive disorders notes.

RULE **Read vascular cognitive decline by event, anatomy and tempo.** Ask whether there was one event, recurrent events or no obvious clinical stroke; whether the damage is cortical, strategically focal or subcortical; and whether the course is stepwise with plateaus or gradually progressive.

6.1 A substrate-based map of vascular cognitive decline

The most useful compact classification recognises **three canonical vascular cognitive impairment subtypes**: multi-infarct, strategic-infarct and small-vessel forms. This arrangement is clinically powerful because it links history to anatomy. Multi-infarct disease reflects cumulative burdens, strategic-infarct disease reflects the disproportionate effect of one critical lesion, and small-vessel disease reflects lacunes or ischaemic white-matter injury. Small-vessel disease is also described as subcortical vascular cognitive impairment, and some authors use Binswanger disease for the white-matter form.

An expanded O'Brien classification places these canonical forms alongside hypoperfusion dementia, haemorrhagic dementia, Alzheimer disease with cerebrovascular disease, and hereditary vascular dementia represented by CADASIL. The expanded taxonomy is not a rival to the compact scheme. The compact scheme is a bedside map of common lesion patterns; the expanded scheme reminds the examiner that vascular cognitive syndromes also arise through reduced perfusion, haemorrhage, mixed pathology and inherited vessel disease.

PATTERN	EVENT RELATION	TYPICAL TEMPO	DOMINANT CLINICAL-ANATOMICAL CLUE
Multi-infarct	Repeated large cortical or subcortical infarction or haemorrhage	Stepwise decline with intervening plateaus	Focal deficits such as hemiplegia, aphasia or apraxia
Strategic-infarct	A single stroke may suffice	New impairment follows the critical event	A lesion in a region essential for cognition
Small-vessel	An obvious clinical stroke may never occur	Usually gradual progression	Hyperreflexia, parkinsonism or gait disturbance
Inherited small-vessel	Recurrent infarcts may occur without important conventional vascular risk factors	Progressive course after recurrent events	Family pattern plus characteristic small-vessel syndrome

The categories may overlap in an individual brain. A person with diffuse small-vessel injury can also sustain a strategic infarct, while repeated lacunes can generate either an apparently stepwise or a smoother decline. Classification should therefore identify the dominant mechanism and acknowledge additional burdens rather than force every patient into a pathologically pure box.

6.2 Multi-infarct disease: the step-and-plateau comparator

Multi-infarct dementia provides the classic comparator for every other vascular course. Each fresh large infarction or haemorrhage adds a new cognitive burden, producing a recognisable step down. Between events, the defining feature is relative clinical stability: there is a plateau without continuing between-step deterioration. The history must therefore describe not merely "progressive decline" but the shape of the graph. Ask what changed abruptly, which new neurological deficit appeared with it, and whether function then stabilised.

The neurological examination makes the history coherent. Prior lesions may leave hemiplegia, aphasia or apraxia, with the deficit corresponding to lesion location. Successive events can create a patchwork profile because each stroke injures a different network. The cognitive burden is cumulative, but its components need not worsen in parallel: language may fall after one event, praxis after another, and motor dependence after a later lesion.

A stepwise history should not be accepted from a vague account of "good days and bad days." Ordinary fluctuation, fatigue, intercurrent illness and environmental demands can produce variable performance without new infarction. Conversely, a patient with repeated small deep infarcts may appear to decline gradually when individual steps are clinically subtle. The clinician reconstructs dated losses, associated focal signs and plateaus from collateral history and records rather than inferring the pattern from the word "stroke."

TRAP **Do not equate all vascular cognitive decline with a stepwise course.** Step-and-plateau decline characterises the classic multi-infarct pattern, whereas small-vessel disease is often gradual and may occur without an obvious clinical stroke.

6.3 Strategic infarct: when one small lesion has a large cognitive effect

A **strategic-infarct vascular cognitive syndrome** occurs when a single stroke damages a site critical for cognitive functioning. Lesion volume alone is therefore an inadequate guide. A comparatively circumscribed infarct can disrupt a relay, hub or connecting pathway on which several cognitive operations depend. The examination answer should make the causal logic explicit: one event, one strategically placed lesion, and a cognitive change anatomically concordant with that lesion.

Classic locations include the thalamus, mesial frontal lobe, caudate and mesial temporal lobe. Dominant-hemisphere language and association systems are represented by the left angular gyrus, while a compact disconnection lesion may involve the genu of the left internal capsule. A second source broadens the high-yield pattern to dominant thalamic and angular-gyrus infarcts, bilateral infarcts, deep frontal infarcts and left-hemisphere infarcts. These are classic examples, not an exhaustive atlas; infarcts elsewhere can also be strategic.

Clinical reasoning proceeds from network to deficit. Thalamic or deep frontal injury may disturb initiation, attention, processing and memory organisation; mesial temporal injury may make memory impairment prominent; angular-gyrus injury may produce a cognitively disabling dominant-hemisphere association syndrome. The exact bedside profile depends on the damaged network and coexisting lesions. What matters diagnostically is the disproportion between a limited lesion burden and a major functional cognitive consequence.

MNEMONIC

S-T-E-P

Stroke relation: single, repeated or clinically silent; **T**empo: abrupt, stepwise or gradual; **E**xamination: focal cortical signs or subcortical motor findings; **P**attern of lesions: strategic hub, multiple territories, lacunes or diffuse white matter.

A strategic lesion should be distinguished from a merely incidental infarct. Temporal proximity helps, but anatomical concordance is the stronger argument. The clinician should describe the new cognitive domain failure, show how it altered independent function, relate it to the damaged region and check for competing diffuse or degenerative burdens. The general diagnostic thresholds for vascular cognitive impairment belong to the dedicated vascular cognitive impairment section.

6.4 Small-vessel and lacunar cognitive disease

The absence of a dramatic stroke does not exclude a vascular substrate. In **small-vessel vascular cognitive impairment**, decline is typically gradual and the patient may never report an event that caused an immediate, obvious change in neurological function. The substrate may be lacunar strokes or ischaemic demyelination. This is why an interview centred only on remembered hemiplegic or aphasic attacks misses an important vascular phenotype.

The bedside phenotype is usually subcortical. Examination may reveal hyperreflexia, parkinsonism and gait disturbance. These findings direct attention away from a purely amnestic

account and toward distributed frontal-subcortical disconnection. Cognitive slowing, reduced mental flexibility, impaired initiation and executive inefficiency may dominate functional difficulty even when the patient or family describes the problem simply as “memory loss.” The clinician must translate that complaint into domains rather than assume that storage failure is primary.

Lacunar dementia has more than one tempo. It may follow a classic stepwise course accompanied by recognisable lacunar syndromes such as pure motor stroke or pure sensory stroke. Alternatively, cumulative lacunes may produce a more gradual decline. Relevant territories include the thalamus, caudate, anterior limb of the internal capsule and orbitofrontal region. Repeated disruption across these nodes can progressively degrade the circuits that organise goal-directed behaviour.

Binswanger disease refers here to widespread white-matter disease producing progressive dementia. The characteristic cognitive emphasis is slowed thinking and executive deficit, with amnestic impairment generally less prominent. This relative pattern is examinable: it does not mean memory is normal, only that executive slowing is more salient than the memory-led stereotype of dementia.

Location alone cannot always separate strategic from cumulative small-vessel disease. The thalamus and caudate appear among strategic sites, while thalamic, caudate and internal-capsule lacunes can contribute cumulatively to lacunar dementia. The distinction is causal configuration. In the strategic pattern, one lesion is sufficient and the cognitive change is closely linked to that event. In lacunar disease, several small lesions or more diffuse ischaemic injury progressively degrade the circuit. A good formulation therefore states the site and the burden: “strategic thalamic infarct” and “cumulative thalamostriatal lacunar disease” make different claims even though their anatomical vocabulary overlaps.

This distinction also disciplines radiological interpretation. A scan report listing lacunes or white-matter disease does not by itself describe the cognitive syndrome. Imaging must be reconciled with tempo, domains affected, focal signs and functional change. Equally, a small reported lesion should not be dismissed when its location plausibly explains a major new cognitive deficit.

> **PITFALL**
>
> Do not dismiss gait change, parkinsonism or brisk reflexes as unrelated age-associated findings when cognitive slowing is also developing. Their convergence may be the clinical signal that the cognitive syndrome is subcortical and vascular.

6.5 Post-stroke dementia: frequency, delay and determinants

Post-stroke dementia is an outcome category, not a synonym for multi-infarct dementia. Across studies, differences in population, timing and definitions produce a wide range. One source reports 15-30% at 3 months after stroke and 20-25% developing delayed dementia. Another gives a broader **6-32% developing dementia of some type after stroke**. These estimates are broadly compatible when their distinct sampling and definitions are preserved.

15-30%	20-25%	47%	48%
3-MONTH DEMENTIA	DELAYED DEMENTIA	MCI AFTER LACUNAR CVA	COGNITIVE IMPAIRMENT IN CADASIL STUDY

Risk remains elevated long after the index event. A longitudinal study found a **twofold increased risk of dementia over 10 years** after stroke. Cognitive morbidity also extends below the dementia threshold: in a study limited to lacunar ischaemic cerebrovascular accidents without large-vessel stroke, 47% had mild cognitive impairment, which was more common than physical disability. A minor motor outcome must therefore not reassure the clinician that cognition is intact.

The major post-stroke determinants are multi-infarct, recurrent and strategic strokes. Silent infarcts and cerebral white-matter lesions also increase risk. The probability is consequently shaped by cumulative tissue burden, recurrence and location, not simply by whether the presenting stroke seemed large. One pathological study identified lacunar infarcts as the most common stroke type seen in vascular dementia. Additional substrates included cerebral amyloid angiopathy in about 10% of pure vascular dementia and hippocampal sclerosis in almost 30% of vascular dementia, the latter possibly reflecting particular ischaemic vulnerability of hippocampal neurons.

These data justify longitudinal surveillance. Baseline cognitive and functional change should be reconstructed after the acute event, then reviewed when independence, adherence, rehabilitation or caregiver burden changes. Screening alone is not the diagnosis, and a delayed syndrome should not be rejected merely because cognition appeared preserved during the initial admission.

6.6 Risk architecture and prevention reasoning

Major risk factors for mild or major vascular neurocognitive disorder are shared with cerebrovascular disease and stroke. They include hypertension, diabetes, smoking, obesity, high cholesterol, high homocysteine, other atherosclerotic or arteriolosclerotic risks, atrial fibrillation and other states that increase cerebral emboli. Cerebral amyloid angiopathy, through haemorrhage, is also important. The list matters because much of the causal burden is potentially modifiable even when established infarct damage is not reversible.

Prevention reasoning should be layered rather than reduced to a generic instruction to “control vascular risk”:

- Identify the mechanism of the index event and the conditions that make recurrence likely.
- Address modifiable vascular and metabolic risks through coordinated medical care and sustained behavioural change.
- Review smoking, weight, metabolic disease and adherence without treating cognitive impairment as a moral failure.
- Look for silent cumulative burden when gait, executive function or behaviour changes despite no remembered new stroke.

- Use family history and absence of expected conventional risks to decide when an inherited arteriopathy requires consideration.

The psychiatrist's role is not to duplicate stroke medicine. It is to recognise how cognitive impairment undermines risk reduction: disorganisation affects medication use, slowed processing limits education, apathy-like behaviour can reduce activity, and caregiver strain can destabilise routines. Instructions should be simplified, collateral support enlisted with consent, and capacity assessed for the decision at hand. Current stroke guidance should direct mechanism-specific prevention.

> **GOOD TO REMEMBER**
> Ask separately about cognitive change, gait and executive function after even a physically mild lacunar event. Physical recovery and cognitive recovery are not interchangeable outcomes.

6.7 CADASIL: inherited subcortical arteriopathy

CADASIL expands to cerebral autosomal-dominant arteriopathy with subcortical infarcts and leukoencephalopathy. It is a subcortical small-vessel syndrome characterised by lacunar strokes, migraine and dementia. Most vascular cognitive impairment is sporadic, so the diagnosis becomes especially important when the age, recurrence pattern, migraine history, conventional risk profile and pedigree do not fit ordinary acquired small-vessel disease.

The genetic lesion involves **NOTCH3 on chromosome 19**, with linkage described at 19q12 and autosomal-dominant inheritance. The gene is expressed in vascular smooth-muscle cells and contributes to cell proliferation and differentiation. The resulting arteriopathy affects penetrating vessels: infarcts arise from thickening and fibrosis of small and medium penetrating arterial walls.

The clinical sequence is high yield. The syndrome first appears after the third decade of life, often through repeated migraine with aura. It then progresses over one decade with recurrent cerebral infarcts despite the absence of important conventional cerebrovascular risk factors, culminating in dementia, usually with pseudobulbar palsy. This is a natural-history pattern, not a rigid schedule for every affected person.

Epidemiologically, it is the most common hereditary stroke syndrome, with prevalence **up to 1 in 20,000 adults**. One study found cognitive impairment in 48% of affected individuals. Yet inherited disease remains a minority of vascular dementia overall: familial forms account for **less than 10% of vascular dementia**. Thus, a strong inherited pattern should sharpen investigation, but indiscriminate genetic attribution is poor reasoning.

6.8 Inherited vascular differential and final synthesis

Familial vascular dementia is uncommon but heterogeneous. CADASIL and hereditary cerebral haemorrhage with amyloidosis-Dutch type are among the better studied syndromes. The differential should be organised by vessel pathology, associated organ disease and inheritance rather than memorised as an undifferentiated gene list.

SYNDROME	INHERITANCE OR GENETIC BASIS	VASCULAR-SYSTEM CLUE
Hereditary cerebral haemorrhage with amyloidosis-Dutch type	Autosomal dominant; APP gene	Amyloid in meningocortical microvasculature with recurrent lobar haemorrhage
Familial British dementia with amyloid angiopathy	Autosomal dominant; BRI2, also called ITM2B, on chromosome 13	Amyloid angiopathy within a familial dementia syndrome
HERNS	Autosomal dominant	Endotheliopathy with retinopathy, nephropathy and stroke
Fabry disease	X-linked recessive alpha-galactosidase A deficiency	Lysosomal storage disease with vascular involvement
MELAS	Mitochondrial transfer-RNA mutation; maternal inheritance	Myopathy, encephalopathy, lactic acidosis and stroke-like episodes

Genetic mechanisms also caution against an absolute vascular-degenerative divide. APOE variants influence cerebrovascular risk, while presenilins interact with NOTCH3. This suggests a **mechanistic overlap between Alzheimer and vascular disease** that may contribute to their coexistence in mixed dementia. The inference is overlap, not equivalence: a vascular history does not exclude additional pathology, and a degenerative phenotype does not make vascular lesions irrelevant.

A complete answer therefore begins with substrate, not with a catch-all label. Contrast multi-infarct step-and-plateau decline with the gradual, sometimes clinically silent small-vessel course. Explain why a single infarct can be strategic, name the critical regions with correct laterality, and show that lacunar and white-matter disease favour a slowed executive-subcortical profile. Then quantify post-stroke cognitive risk, identify cumulative and modifiable determinants, and recognise the inherited pattern of migraine with aura, recurrent lacunar events, sparse conventional risk factors and autosomal-dominant transmission.

> One critical infarct, many cumulative infarcts and diffuse small-vessel injury are different routes to vascular cognitive decline; CADASIL is the inherited small-vessel prototype.

Traumatic brain injury and its psychiatric sequelae

7 · Traumatic brain injury: classification and the primary versus secondary injury model

The high-yield distinction is between grading the injury and explaining how it evolves. Traumatic brain injury is not described adequately by a scan label or by whether the patient recalls being unconscious. A defensible formulation uses converging clinical anchors to grade severity, describes the mechanical pattern of damage, and then separates damage produced at impact from processes set in motion afterward. That sequence matters because the severity label, anatomical lesion and biological phase answer different questions.

IN INDIA

Take an explicit head-injury history whenever a psychiatric presentation follows a road crash, even one that seemed minor, since road traffic injury is a common mechanism of traumatic brain injury.

The marks usually sit at the junctions: how the **Glasgow Coma Scale**, loss of consciousness and post-traumatic amnesia are integrated; why the best available score after early stabilisation is more useful than an isolated immediate score; how mild injury can coexist with a structural lesion; and why secondary injury is potentially dynamic even though the initiating trauma has ended. For psychiatrists, this framework also prevents later behavioural or cognitive change from being read as a simple linear consequence of the initial score. The accompanying figure brings the severity bands and the injury cascade together.

13-15	9-12	3-8	24 hours
MILD GCS BAND	MODERATE GCS BAND	SEVERE GCS BAND	BEST-SCORE WINDOW

7.1 Classification is multidimensional

Severity classification begins with the acute physiological disturbance, not with later symptom burden. The conventional scheme combines level of consciousness, duration of unconsciousness and duration of post-traumatic amnesia. These variables overlap but are not interchangeable. A depressed consciousness score is a sampled bedside state, while loss of consciousness and post-traumatic amnesia supply duration anchors. This fragment uses amnesia only to grade severity. Its full phenomenology, assessment and rehabilitation implications belong in the section on cognitive impairment, post-traumatic amnesia and rehabilitation after traumatic brain injury.

An anatomical axis describes the pattern of damage. Trauma may produce local damage beneath an impact, damage opposite the impact, contusions at vulnerable surfaces, or distributed white-matter injury from rotational and shearing forces. A biological-phase axis distinguishes immediate mechanical disruption from delayed physiological and cellular consequences. Thus, a complete description should avoid collapsing severity, morphology and phase into a single adjective.

RULE **Classify severity with the acute clinical anchors, then describe anatomy and phase separately.** A mild, moderate or severe label is not a substitute for naming the lesion pattern, and an anatomical lesion does not retrospectively rewrite every acute severity parameter.

This separation improves clinical reasoning. The initial classification supports communication and broad prognostic framing, whereas morphology explains which tissue was exposed to mechanical stress and the phase model explains why injury can evolve after the accident. Later psychiatric syndromes require their own assessment. Mood, anxiety, trauma-related, psychotic, personality and behavioural sequelae are taught in their dedicated traumatic-brain-injury sections rather than being inferred directly from a severity category.

A retrospective psychiatric assessment should preserve the quality of the evidence. Contemporary emergency records, serial neurological observations, imaging reports and

collateral descriptions may not agree perfectly, and later recall is often least secure around the period being classified. State which anchors were observed, which were reconstructed and which remain unknown. If the score was obtained when verbal performance could not be tested, preserve that limitation. If loss-of-consciousness duration is only estimated, do not convert it into false precision. This discipline is more informative than assigning a confident category from one ambiguous entry, and it leaves later reviewers able to distinguish evidence from inference.

7.2 Glasgow Coma Scale: structure, timing and limitations

The Kaufman's Clinical Neurology for Psychiatrists account describes the scale as measuring three readily observable neurological functions: eye opening, best verbal response and best motor response. Their combined total extends from **3 to 15**, with lower totals indicating poorer neurological function. The score is compact and reproducible enough to support handover, but it should be recorded as a clinical observation rather than treated as an explanation of the injury.

The severity bands are 13 to 15 for mild injury, 9 to 12 for moderate injury, and 3 to 8 for severe injury. Timing matters. The best score in the first **24 hours after injury** is more useful than the score obtained immediately after injury. The reason is interpretive: a single earliest observation can reflect the unstable circumstances surrounding trauma and resuscitation, while the best documented early score more usefully represents the patient's neurological capacity during that period.

The scale also has important limits. It is not readily applicable after cerebral hypoxia. Intubation prevents a verbal response, so a total that is quoted without describing this constraint may look more precise than it is. The correct response is not to discard the assessment but to record what could actually be tested, identify the reason a component is unavailable, and interpret the total in clinical context.

> **PITFALL**
>
> Do not copy an undocumented Glasgow Coma Scale score from a referral and let it dominate the formulation. Establish when it was obtained, whether verbal testing was possible, whether cerebral hypoxia limits interpretation, and whether a better score was documented during the early post-injury window.

7.3 The integrated mild, moderate and severe bands

The classic classification places loss of consciousness, post-traumatic amnesia and the Glasgow Coma Scale side by side. The table should be read across each row. A patient is not graded intelligently by searching for whichever isolated value produces the most dramatic label; discordant anchors should be documented, explained where possible and interpreted with the rest of the injury evidence.

SEVERITY	DURATION OF UNCONSCIOUSNESS	GLASGOW COMA SCALE	POST-TRAUMATIC AMNESIA	INTERPRETIVE POINT
Mild	Less than 30 minutes	13-15	Less than 24 hours	The label grades the acute injury; it does not mean that every later complaint will be minor
Moderate	30 minutes to 24 hours	9-12	24 hours to one week	Converging anchors strengthen the classification; disagreement requires clinical interpretation
Severe	More than 24 hours	3-8	More than one week	This is the severe row in the classic three-band scheme, not the alternative definition discussed below

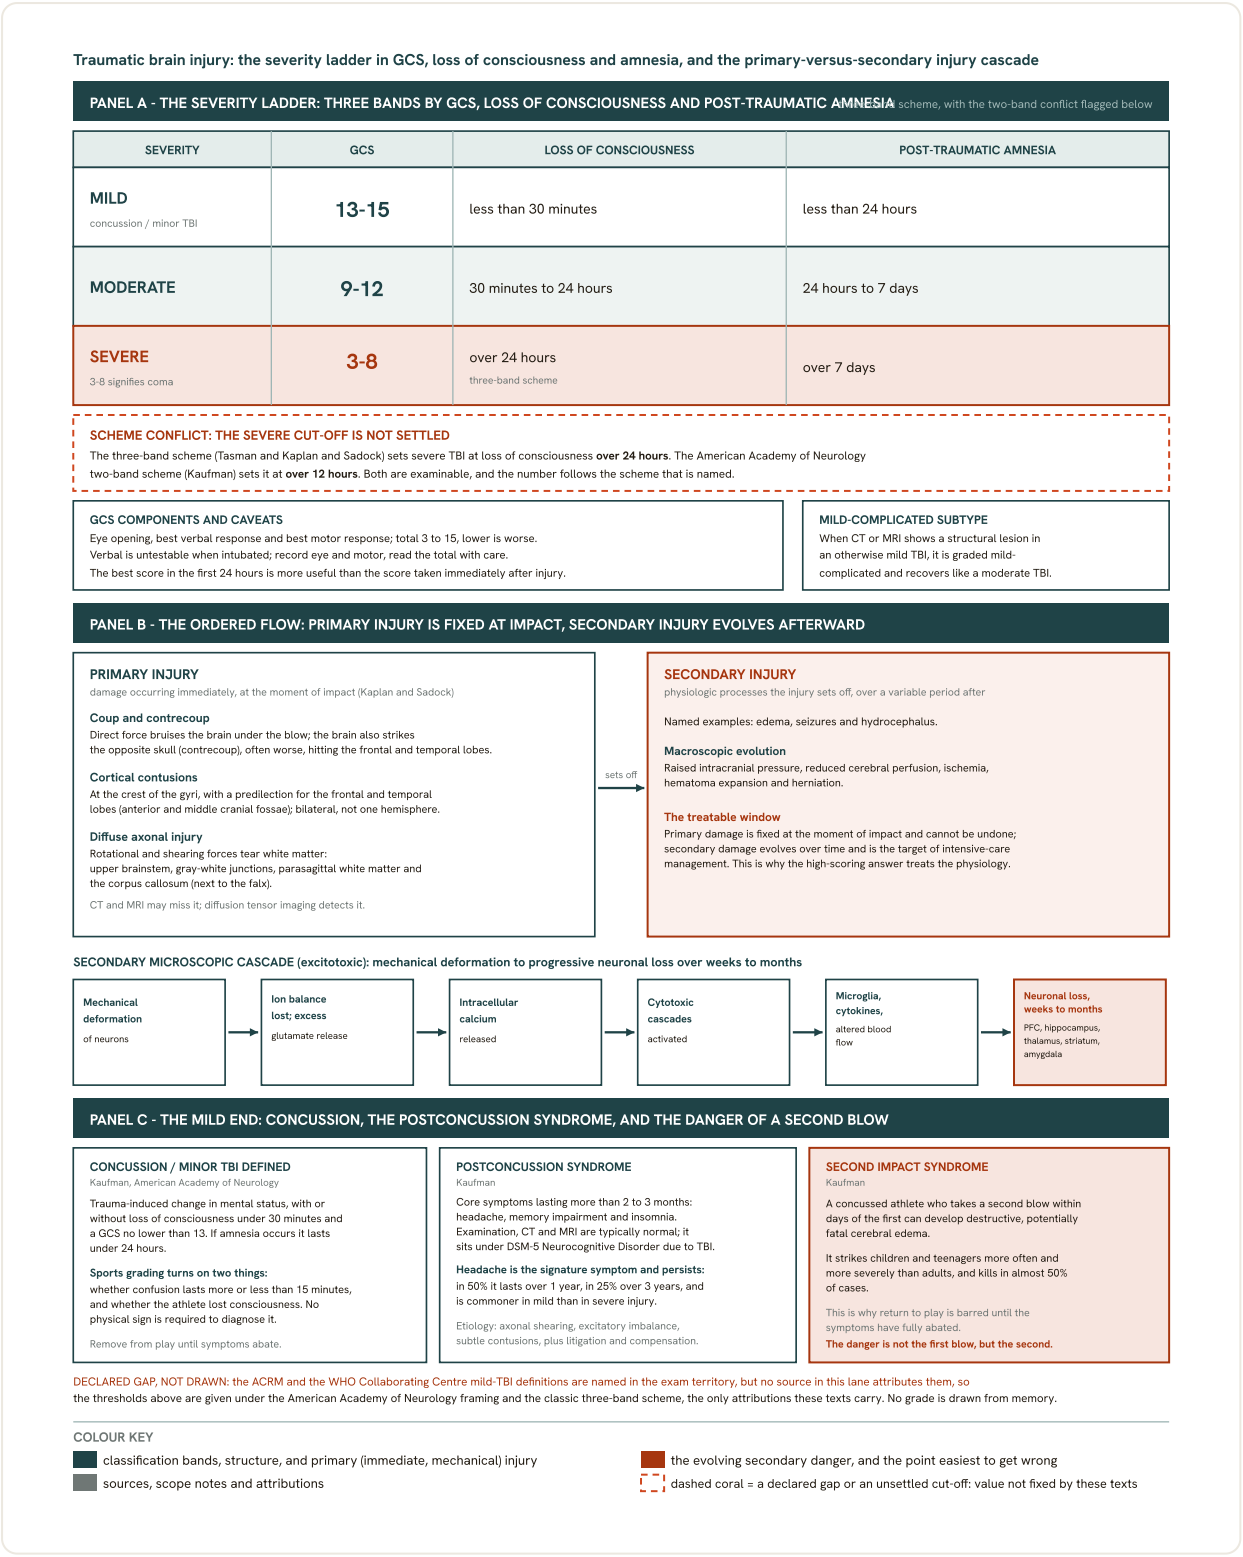


Fig 34.3 **Traumatic brain injury: the severity ladder and the primary-versus-secondary cascade.** The three-band scheme grades severity by GCS (13-15, 9-12, 3-8) with loss of consciousness under 30 minutes, up to 24 hours, or over 24 hours, while the American Academy of Neurology two-band scheme calls loss of consciousness over 12 hours severe, so the cut-off follows the scheme named. Primary injury (coup and contrecoup, contusions, diffuse axonal injury) is fixed at impact; secondary injury (edema, raised pressure, ischemia and the excitotoxic cascade) evolves afterward and is the treatable window. (Kaplan and Sadock, CTP 2024; Kaufman's Clinical Neurology for Psychiatrists; Tasman Psychiatry.)

Each anchor has a distinct vulnerability. The consciousness interval may be uncertain when the event was unwitnessed or documentation was delayed. Amnesia is partly reconstructed from the patient's later account and collateral history. The Glasgow Coma Scale is a sampled examination

affected by testability and timing. Using the anchors together therefore adds structure without pretending that uncertainty has disappeared.

When the anchors disagree, first check whether they refer to the same clinical period. An immediate score and a later best score are not rival measurements of an unchanging state. Next ask whether a component was untestable and whether the time intervals came from direct observation or retrospective estimation. Then report the discordance rather than manufacturing a neat answer. The classification is a disciplined summary of evidence, not a mathematical rule that makes missing or incompatible observations vanish. Imaging can add morphology, as in mild-complicated injury, but it does not erase the historical consciousness and amnesia data.

MNEMONIC

CAP

Consciousness duration, **A**mnesia duration and **P**erformance on the Glasgow Coma Scale form the severity spine. Then add anatomy and phase. CAP is a retrieval aid, not a rule for forcing discordant data into a single category.

7.4 Mild injury, concussion and mild-complicated injury

Under the American Academy of Neurology framing reported by Kaufman, **mild traumatic brain injury** is a trauma-induced alteration in mental status that may occur with or without loss of consciousness. The alteration typically consists of confusion and amnesia lasting seconds to minutes, without focal neurological signs such as hemiparesis, cranial-nerve abnormality or incoordination. Physicians may use **concussion** for mild traumatic brain injury. This conceptual definition complements the numerical mild band without making loss of consciousness compulsory.

The term mild describes initial severity. It must not be used as shorthand for no brain disturbance, no need for observation, or guaranteed absence of later symptoms. Conversely, later nonspecific symptoms do not by themselves reconstruct the acute injury as a concussion. The full post-concussion syndrome, its persistent symptom pattern and its management belong in the section on mild traumatic brain injury, concussion and post-concussion syndrome.

Imaging supplies an additional axis. Some patients who otherwise meet mild traumatic brain injury criteria have a structural lesion demonstrated by CT or MRI. This is **mild-complicated traumatic brain injury**. Its recovery trajectory resembles that of moderate injury, which leads some classifications to place the subtype with moderate traumatic brain injury. The crucial reasoning is that a mild clinical band and a positive structural study are not mutually exclusive. One describes the acute physiological severity anchors; the other demonstrates morphology.

GOOD TO REMEMBER

Write the formulation in layers: acute severity band, available Glasgow Coma Scale timing and limitations, duration anchors, and imaging status. “Mild-complicated” is precisely the situation in which reducing the case to either “mild” or “scan positive” loses useful information.

7.5 Severe injury: the threshold conflict

There is no honest way to print one loss-of-consciousness threshold for severe traumatic brain injury without naming the scheme. In the classic three-band classification, severe injury includes unconsciousness for **more than 24 hours**, post-traumatic amnesia for more than 7 days, and a Glasgow Coma Scale score of 3 to 8. These are the severe values in the integrated mild, moderate and severe table.

A different two-category neurological framing reported by Kaufman defines **severe traumatic brain injury** as prolonged post-traumatic loss of consciousness for **more than 12 hours** together with radiographic evidence of injury to the brain, skull or intracranial blood vessels. That definition changes more than the threshold. It also includes an imaging requirement and belongs to a different classificatory frame.

TRAP **Do not merge "more than 12 hours" and "more than 24 hours" into one supposedly universal severe-TBI cut-off.** The former belongs to a two-category definition that also requires radiographic injury; the latter belongs to the classic three-band loss-of-consciousness table. State the scheme, then state its threshold.

This conflict is an example of why criteria must travel with their qualifiers. If an examination asks for the mild, moderate and severe table, give the integrated three-band values. If discussing the alternative neurological definition, identify its paired radiographic requirement. Do not average the thresholds, silently choose one, or imply that the sources describe the same construct with casual wording differences.

7.6 Primary injury: impact, coup, contrecoup and contusion

Primary brain injury is the tissue damage occurring immediately at the time of trauma. It is the direct mechanical event. The pattern depends on how the head and brain move, where force is transmitted and which tissue interfaces bear deformation. This category includes local impact damage as well as distributed injury from rotational and tensile-shearing forces. Calling it primary is temporal and mechanistic; it does not mean clinically unimportant or necessarily focal.

A direct blow can disrupt brain tissue beneath the site of impact, producing a **coup injury**. The same trauma can propel the brain against the opposite internal surface of the skull, producing a **contrecoup injury**. Contrecoup damage frequently exceeds coup damage. The frontal and temporal lobes are especially vulnerable because their anterior surfaces meet the sharp contours of the anterior and middle cranial fossae.

Traumatic cerebral contusions may occur at the crests of gyri and show a predilection for frontal and temporal regions housed in the anterior and middle cranial fossae. This anatomical pattern helps explain why apparently remote impact can coexist with damage at the brain's basal and anterior surfaces. It also cautions against assuming that the site of the external blow identifies the only injured region.

- **At impact:** ask where force entered, but also consider movement of the brain within the skull.

- **Across the skull:** distinguish tissue beneath the blow from damage at the opposite internal surface.
- **At vulnerable surfaces:** remember the frontal-temporal predilection of contusions.

7.7 Traumatic axonal injury: distributed mechanical damage

Traumatic axonal injury, also called diffuse axonal injury, consists of multifocal white-matter damage produced by rotational and tensile-shearing forces. “Diffuse” should not be read as uniform damage to every part of the brain. The defining idea is distributed, multifocal axonal injury created by mechanical deformation, often involving bilateral or midline structures rather than a single cortical impact site.

Commonly injured regions include the upper brainstem and diencephalic structures, grey-white matter junctions, parasagittal cerebral white matter and the corpus callosum. The corpus callosum is particularly sensitive because it lies near the falx. This distribution follows the physics of tissues moving differently at interfaces and around relatively fixed structures. It is therefore conceptually distinct from a coup contusion, even though the mechanisms can coexist in the same traumatic event.

The distinction between focal impact injury and distributed axonal injury is useful to psychiatrists because neither the external wound nor a single focal lesion necessarily captures the functional network disturbance. It does not license a retrospective microscopic diagnosis from symptoms alone. The mechanism must remain an anatomically plausible component of the injury formulation, supported by the available neurological assessment and investigations.

A concise mechanical classification can therefore be retained as follows:

- **Local impact pattern:** coup injury, contrecoup injury and surface contusion.
- **Distributed shearing pattern:** multifocal axonal injury at characteristic white-matter and midline regions.

7.8 Secondary injury: physiological and cellular amplification

The accident ends; the biology does not necessarily stop. Secondary brain injury refers to physiological processes initiated by trauma that unfold over a variable period afterward. Examples include oedema, seizures and hydrocephalus. The wider cascade includes raised intracranial pressure, reduced cerebral perfusion, ischaemia, expansion of a haematoma and herniation. These are consequences of the original insult, but they are not identical to the instantaneous mechanical damage.

At cellular level, mechanical deformation disrupts ion homeostasis, drives excessive release of glutamate and other neurotransmitters, releases intracellular calcium and activates cytotoxic cascades. Delayed processes also include microglial activation, inflammatory cytokine release and altered cerebral blood flow. These mechanisms link an immediate physical event to continuing metabolic stress, inflammation, impaired perfusion and further neuronal injury.

Damage can continue over weeks to months in vulnerable regions including prefrontal cortex, hippocampus, thalamus, striatum, amygdala and forebrain nuclei. This prolonged biological

course does not mean that every later symptom proves ongoing tissue destruction. It means that a purely static model of trauma is incomplete. Clinical change after injury requires reassessment of neurological state, complications, medication and substance effects, sleep, pain, psychiatric syndromes and environmental demands rather than automatic attribution to either the original impact or the secondary cascade.

The primary-secondary distinction also prevents a common conceptual error: treating secondary injury as a second accident. It is an evolving consequence of the first insult. Primary damage is mechanically delivered at impact; secondary damage is mediated through later physiological and cellular processes. Raised intracranial pressure, hydrocephalus and post-injury seizures are named here as parts of that cascade, while their detailed neuropsychiatric assessment belongs in the dedicated section on those complications. Rehabilitation and pharmacological management likewise belong in their respective traumatic-brain-injury sections.

For examination writing, the causal chain should be explicit. Begin with deformation at impact, identify the focal or distributed mechanical pattern, and then show how disturbed perfusion, pressure, excitatory signalling, calcium handling and inflammation can amplify injury afterward. Do not list these as unrelated complications. Their teaching value lies in the sequence from mechanical force to physiological instability and cellular toxicity. Equally, do not imply that every secondary process occurs in every patient. The model supplies a framework for surveillance and interpretation; the individual formulation depends on observed clinical and investigative evidence.

> Grade traumatic brain injury with Glasgow Coma Scale, loss of consciousness and post-traumatic amnesia, then separate immediate mechanical primary injury from the evolving physiological and cellular processes of secondary injury.

8 · Mild traumatic brain injury, concussion and the post-concussion syndrome

Mild injury does not mean trivial injury, and a normal scan does not mean an imaginary syndrome. **Mild traumatic brain injury**, often discussed clinically as concussion, is defined by a trauma-related change in mental status within specified severity limits. The examination marks lie in separating what qualifies as mild injury, what makes persistent symptoms clinically coherent, and why symptom burden may not track conventional injury-severity markers.

> **IN INDIA**
>
> Give the patient and family clear written advice on rest, staged return and the danger of a second blow, and do not assume that trained sideline supervision is in place.

The psychiatrist commonly enters after the emergency narrative has faded. The patient may now be worried by headache, poor memory or disturbed sleep; the family may be uncertain whether recovery is occurring; and routine examination or imaging may be unrevealing. A sound answer neither dismisses symptoms nor converts every complaint into proof of structural damage. It

reconstructs the injury, defines the time course, recognises the characteristic symptom pattern, searches for competing explanations, and plans review according to function and risk.

50%	25%	Usually normal	Almost 50%
HEADACHE BEYOND 1 YEAR	HEADACHE BEYOND 3 YEARS	EXAMINATION AND IMAGING	SECOND-IMPACT MORTALITY

8.1 Operational definition of mild traumatic brain injury

The operational definition begins with **trauma-induced alteration in mental status**. Loss of consciousness may occur, but it is not mandatory. When consciousness is impaired, the duration remains **less than 30 minutes**. The **Glasgow Coma Scale** score is **no lower than 13**. If amnesia occurs, it lasts **less than 24 hours**. These limits identify the mild end of traumatic brain injury; they do not predict that every later symptom will be mild.

This definition is easy to distort in an examination answer. Concussion is not synonymous with being knocked unconscious. The positive requirement is an injury-related alteration in mental status; impaired consciousness is an optional feature constrained by duration. Likewise, the amnesia limit is part of severity placement, not an invitation to teach the broader construct of post-traumatic amnesia here. The full severity classification and the primary versus secondary injury model belong in the section on traumatic brain injury classification and injury mechanisms; the full post-traumatic amnesia construct belongs with cognitive impairment and rehabilitation after traumatic brain injury.

RULE **Define concussion by altered mental status after trauma, not by mandatory loss of consciousness.** Then state the consciousness, Glasgow Coma Scale and amnesia limits without turning a mild-severity label into a prognosis.

The operational band performs a classificatory job. It gives clinicians a common language for the initial injury and helps organise acute communication. It cannot, by itself, settle whether a later headache is genuine, whether cognitive inefficiency is functionally important, or how long recovery will take. That distinction between classification and outcome is the conceptual hinge of this topic.

8.2 Clinical assessment after the injury

Assessment should begin with a chronology rather than a symptom inventory. Establish the injury event, the immediate alteration in mental status, any impairment of consciousness, any amnesia, and the evolution of complaints since then. The clinician is trying to connect a present syndrome to a plausible index injury while remaining open to complications and alternative causes. A vague history of “head injury sometime earlier” is insufficient; an equally vague declaration that the scan was normal is not a rebuttal.

A focused assessment can be organised around a short clinical sequence:

- **Reconstruct the index event:** clarify the mental-status change and document the available severity anchors without inferring details that were never observed.

- **Describe the present complaints:** ask about headache, memory difficulty and insomnia, then translate each complaint into its effect on work, study, self-care, relationships and ordinary routines.
- **Establish the trajectory:** determine what improved, what persisted and what newly appeared, because a changing pattern deserves fresh clinical explanation rather than automatic attribution to concussion.
- **Seek corroboration:** use records and collateral history where available, especially when the patient cannot confidently reconstruct the immediate post-injury period.

The neurological and mental-state examinations answer different questions. The neurological examination searches for evidence that would redirect the case beyond an uncomplicated mild-injury formulation. The mental-state examination characterises distress, attention during the interview, confidence in memory, sleep-related preoccupation, mood and the meaning the patient assigns to symptoms. Neither examination should be used as a credibility test. Observable performance and subjective difficulty may diverge without proving deception or absence of impairment.

> **GOOD TO REMEMBER**
>
> Document function alongside symptoms. “Memory is poor” is diagnostically weak; examples of lost task efficiency, repeated errors, abandoned routines or reduced tolerance for ordinary cognitive load make the complaint assessable over time without pretending that bedside observation alone measures its cause.

If the course is atypical, worsening or clinically unstable, do not protect the original label. Reassess the patient and widen the medical differential through the appropriate emergency or neurological pathway. Detailed imaging decision rules and red-flag lists should be applied from current head-injury guidance and the clinical setting rather than reconstructed from memory.

8.3 Post-concussion syndrome: the persistent clinical pattern

Kaufman's Clinical Neurology for Psychiatrists notes that **post-concussion syndrome** lacks a strict definition. Neurologists base the diagnosis on persistence of one or more core symptoms, especially **headache, memory impairment and insomnia**, for **more than 2 to 3 months**. The formulation therefore combines a preceding mild traumatic brain injury, a recognisable symptom pattern and persistence beyond the expected early period. It is not established by duration alone.

The term is useful when it organises follow-up, but it can mislead if treated as a unitary lesion with a single biomarker. The symptoms are subjective and nonspecific, and routine assessment often supplies no positive structural marker. This is a clinical diagnosis built from history, symptom pattern, course and exclusion of better explanations, not a radiological diagnosis waiting for a visible lesion. The typical examination and imaging profile is set out in the comparison below.

CLINICAL AXIS	MILD TBI OR CONCUSSION	PERSISTENT POST-CONCUSSION PRESENTATION	INTERPRETIVE CAUTION
Anchor	Trauma-related alteration in mental status within mild-injury limits	A prior concussion followed by continuing complaints	Do not diagnose concussion from later nonspecific symptoms alone
Consciousness	Impairment may be present or absent	Historical feature, not a requirement for current symptom validity	Do not demand a blackout
Core symptom anchor	Immediate mental-status alteration	Headache, memory impairment or insomnia	The symptoms remain nonspecific and require clinical formulation
Examination and imaging	Used to assess the injury and competing pathology	Physical examination, neurological examination, CT and MRI are typically normal	Normal findings neither prove nor disprove the subjective burden
Severity relation	Classifies the initial injury	Symptoms do not correlate with Glasgow Coma Scale score or amnesia duration	Initial severity is not a symptom meter

TRAP **Do not write "normal CT or MRI excludes post-concussion syndrome."** Routine structural studies and examination are typically normal; their principal value is assessing other pathology, not validating or invalidating every persistent complaint.

Diagnostic discipline requires a balanced stance. Accept the patient's report as clinically relevant, define its functional consequences, and still examine for other neurological, psychiatric, sleep-related and contextual explanations. This is not scepticism. It is the ordinary causal work required whenever a common, nonspecific symptom follows a memorable event.

8.4 Post-concussion headache

Post-concussion headache is the signature and most common post-concussion symptom in the clinical account. Its persistence is clinically important, as shown in the accompanying statistics strip. The same account reports that post-concussion headache occurs more often after mild injury than after moderate or severe injury. This counterintuitive pattern is an examination favourite because it blocks the lazy assumption that symptom frequency must rise in a straight line with injury severity.

Headache should be described, not merely named. Establish its relation to the injury, subsequent course, functional effect and associated concerns. Then examine the patient and decide whether the present pattern remains compatible with the working post-concussion formulation or demands renewed investigation. The label "post-concussion" is an aetiological hypothesis. It must not stop the clinician from reconsidering causation when the presentation changes.

The persistence figures also change communication. A clinician who promises rapid disappearance may unintentionally intensify alarm when symptoms continue. Equally, presenting persistence as inevitable encourages passivity and catastrophising. The better message is conditional: prolonged headache is recognised, its duration does not map neatly onto initial severity, and management should remain active, symptom-focused and function-oriented while alternative explanations are reviewed when appropriate.

PITFALL

Do not use injury severity to decide whether headache is believable. The clinical pattern is paradoxical: headache is reported more frequently after mild injury than after more severe injury, and persistent symptoms do not correlate neatly with the usual initial severity parameters.

Specific headache phenotyping, acute investigation and treatment should follow current headache and head-injury guidance and the patient's medical context. In an examination answer, it is safer and more accurate to state the management principles than to invent a familiar dose.

8.5 Why symptoms persist: a biopsychosocial formulation

Proposed neurological explanations include **diffuse axonal shearing**, excitatory neurotransmitter imbalance and subtle cerebral contusions. These mechanisms explain why a conventional structural study may be unrevealing despite a biologically plausible injury. They remain proposed explanations, not routine bedside measurements, and should not be presented as if a normal scan has demonstrated microscopic damage in an individual patient.

Prolonged symptoms are also associated with psychiatric and socioeconomic factors, including litigation, compensation, premorbid intellectual characteristics and premorbid personality abnormalities. The correct frame is **biopsychosocial formulation**: neurological injury, symptom interpretation, emotional response, prior vulnerabilities, environmental demands and social contingencies can interact. This formulation places psychiatric and socioeconomic associations alongside neurological injury rather than beneath it.

This requires careful language. Compensation or litigation is not a synonym for fabrication. Premorbid personality is not a moral defect. Psychiatric contribution does not make a symptom unreal. Each factor is a potential modifier of persistence, distress, disability or help-seeking and must be assessed on its own evidence. A formulation should state what is observed, what is inferred and what remains uncertain.

MNEMONIC

MIND

Mental-status alteration defines the injury; **I**njury mechanisms may be subtle; **N**ormal examination or imaging is common in the persistent syndrome; **D**uration and disability require a biopsychosocial formulation. The mnemonic preserves the logic without claiming a single cause.

The absence of correlation with estimated impact force, Glasgow Coma Scale score or amnesia duration is especially important. It means that conventional severity anchors cannot adjudicate the later symptom experience. Their job was to classify the acute injury. The later task is to explain persistence and functional impact. Confusing these jobs produces either overmedicalisation or dismissal.

8.6 Sports concussion and second-impact risk

Sports concussion grading in the clinical account uses the duration of confusion and the presence or absence of loss of consciousness. Confusion is separated by a 15-minute threshold, while loss of consciousness is recorded as present or absent. Physical symptoms or signs are not required to diagnose concussion in the injured athlete. This again reinforces the central principle that concussion is a disturbance of brain function, not a diagnosis that waits for a visible external sign.

AAN guidance recommends removing a concussed athlete from play immediately; the athlete should not return merely because they insist they feel able to continue. The athlete remains out until symptoms have abated and a licensed professional has evaluated recovery. The purpose is not bureaucratic caution. It is prevention of another injury while the brain remains vulnerable.

Second impact syndrome describes a further blow received within days of the original concussion, followed by destructive and potentially fatal cerebral oedema. Children and teenagers are affected more frequently and more severely than adults. Reported mortality is strikingly high, as shown in the accompanying statistics strip. This is the high-stakes reason why apparent determination, competitive pressure or absence of obvious physical signs must never be allowed to substitute for recovery and professional clearance.

The older grading framework is worth knowing because it appears in textbook and examination language, but bedside action should follow current sport-specific guidance. Take any graduated return-to-play protocol, stage intervals and exercise thresholds directly from that guidance rather than inventing them. The non-negotiable principle is to remove the athlete from play, allow symptoms to resolve, obtain professional assessment and prevent repeat injury during vulnerability.

8.7 Diagnostic formulation and communication

A defensible formulation has layers. Begin with whether the index event meets the mild-TBI definition. Then describe the present symptom pattern and its duration. Record examination and imaging findings without treating normality as a verdict on credibility. Identify functional impairment and contextual modifiers. Finally, state competing explanations and the plan for review. This structure keeps diagnosis provisional enough to remain safe and specific enough to guide care.

The clinician should avoid polarised explanations. “It is all brain damage” overstates what routine assessment can prove. “It is only stress” erases the injury and usually damages trust. A more accurate explanation is that concussion is a genuine functional disturbance after trauma, persistent symptoms can occur despite normal routine investigations, and biological,

psychological and social factors may all influence recovery. This allows uncertainty without abandonment.

Separate symptom, impairment and disability in ordinary language. A symptom is what the patient experiences; impairment is the difficulty demonstrated in a particular function; disability is the restriction that emerges when those difficulties meet the person's real-world demands. The distinctions prevent a single complaint from carrying the entire causal argument. They also make follow-up meaningful because improvement may occur in function before complete symptom disappearance, or symptoms may remain while participation expands.

The broader psychiatric sequelae of traumatic brain injury, including mood, anxiety, trauma-related and psychotic presentations, are taught in the dedicated psychiatric-sequelae section. Personality change, aggression and disinhibition have their own section. Cognitive rehabilitation and chronic traumatic encephalopathy also belong elsewhere, as does pharmacological management in TBI psychiatry. Naming these possibilities is necessary; reproducing their criteria or treatments here would blur ownership and weaken the concussion answer.

A clinical letter should therefore communicate the injury basis, current symptoms, objective findings, functional consequences, uncertainty and safety plan. It should avoid implying malingering from normal imaging, avoid promising a fixed recovery date, and avoid using the mild label to minimise present disability. Precise neutral documentation is particularly important when occupational, educational, sport or compensation decisions depend on it.

8.8 Management principles: LOCUS, FOCUS and MODUS

Management is organised by current risk and function, not by the reassuring adjective “mild”.

LOCUS: emergency or inpatient reassessment is appropriate when the patient is unstable or the course raises concern for pathology beyond an uncomplicated concussion. A stable patient with persistent symptoms can usually be followed in an outpatient pathway with access to medical, neurological, psychiatric and rehabilitation expertise as the formulation requires. Sports settings add a firm removal-from-play boundary and professional clearance before return.

FOCUS: in the acute phase, protect the patient from repeat injury, reassess evolving symptoms and communicate uncertainty clearly. In the short term, track headache, memory difficulty, insomnia and functional tolerance rather than relying on a global statement that the patient is “better”. In the middle phase, the target becomes restoration of ordinary activity without interpreting every fluctuation as fresh damage. In the long term, persistent disability requires renewed formulation, attention to psychiatric and socioeconomic modifiers, and review for alternative explanations.

MODUS: treatment is multimodal. Education should validate the injury while preventing catastrophic interpretations. Psychological work can address fear, symptom monitoring, loss of confidence and avoidance when these maintain disability. Social intervention may include coordinated communication with family, educational or occupational systems and sport personnel. Rehabilitation input should be selected according to demonstrated functional need. Medicines, when considered for individual symptoms, belong to the pharmacological-

management section and must follow current guidance, comorbidity, adverse-effect and interaction review.

Follow-up should compare the patient with their own earlier state. Review the symptom pattern, daily participation, sleep, confidence, distress and any new clinical concern. If the working formulation stops explaining the course, reopen it. Persistent post-concussion symptoms are not a licence for endless passive observation; neither are they a licence for escalating unproven interventions. The goal is safe recovery of function with proportionate investigation and treatment.

> Concussion is defined by trauma-related altered mental status, not obligatory loss of consciousness; persistent symptoms can be real despite normal routine examination and imaging, and severity markers do not measure later symptom burden.

9 · Psychiatric sequelae of traumatic brain injury: depression, anxiety, PTSD and psychosis

After traumatic brain injury, the psychiatric diagnosis must explain a change over time, not merely attach a label to distress after trauma. Depression, anxiety, post-traumatic stress disorder, psychosis and secondary mania can each follow traumatic brain injury, may coexist, and can interfere with recovery in different ways. The marks lie in showing disciplined attribution: establish the pre-injury baseline, reconstruct the injury and subsequent course, define the current syndrome, and test whether another neurological, psychiatric, substance-related or treatment-related explanation fits better.

The apparent contradiction in the literature is itself examinable. Reported frequencies vary with case selection, injury severity, timing and diagnostic method, while some syndromes emerge after a clinically quiet interval. A strong answer therefore distinguishes frequency from individual causation and separates symptom resemblance from syndromal diagnosis. It also recognises that psychiatric assessment continues beyond the early recovery period rather than closing when the acute injury has settled.

11%	12% TO 30%	RR 1.65	29.4%
POOLED GAD ESTIMATE	PTSD IN CIVILIAN COHORTS	SUBSEQUENT SCHIZOPHRENIA	MDD AT 12 MONTHS

9.1 Diagnostic frame: syndrome, chronology and attribution

Begin with linked narratives: what the person was like before injury, what happened around the injury, and how the present psychiatric state developed. Premorbid mood, anxiety, psychotic and substance-related illness may recur after trauma; new symptoms may also be an acquired consequence of brain injury or the psychological meaning of the event and its losses. The task is not to force a false choice between “organic” and “psychological.” It is to formulate how cerebral injury, trauma exposure, disability, disrupted roles and premorbid vulnerability converge in the present syndrome.

Bidirectional vulnerability matters when reading the history. People with premorbid psychiatric disorders have a 70% higher risk of sustaining traumatic brain injury. Consequently, an association between injury and later symptoms does not prove that every disorder began with the injury. Conversely, a premorbid diagnosis does not justify dismissing a genuine post-injury change. Compare baseline phenomenology, severity, functional effect and temporal pattern with the current presentation.

Attribution becomes stronger when the syndrome has a coherent onset after injury, represents a clear change from baseline, fits the observed neurological and psychosocial context, and remains after reasonable mimics have been considered. Collateral history is especially valuable because impaired recall, reduced insight or family reinterpretation can distort the chronology. Records may clarify the injury and early behaviour, while relatives can describe change in ordinary function. Neither source should replace direct examination.

RULE **Diagnose the psychiatric syndrome first, then argue its relation to the traumatic brain injury.** Temporal association supports attribution, but chronology, baseline change, phenomenological coherence and exclusion of better explanations make the causal formulation defensible.

9.2 Depression: frequency, timing and persistence

Post-TBI depressive disorder is common but its reported frequency is strikingly heterogeneous. Across the broad literature, estimates range from **6% to 77%**, reflecting differences in diagnostic criteria, sampling time and injury severity. A narrower source estimate places post-traumatic depression at 25% to 50%. These ranges overlap and should be presented as differently framed estimates, not averaged into a new figure.

Risk appears highest during the first year after injury, and the disorder may persist; one study found that two-thirds of affected patients remained symptomatic for a year or longer. Yet surveillance must not end after that early high-risk period. Among those without depression during the first year, about 25% developed it during the second year. The temporal lesson is practical: assess mood repeatedly during rehabilitation and later functional transitions, including return to work, loss of role, persistent disability and changing family expectations.

The more granular diagnostic estimates make the burden easier to compare. At 12 months postinjury, approximately **29.4% met DSM-IV criteria for major depressive disorder, compared with an annual major-depression rate of 6.7% in the general population**; depressive disorder not otherwise specified and dysthymic disorder were estimated at 19.6% and 1%, respectively. These historical diagnostic categories should not be silently converted into current manual criteria. They show the scale and distribution observed in the cited cohort, while contemporary diagnosis still requires the current classification framework.

Controlled comparison supports an injury-related excess without eliminating the contribution of trauma itself. In a longitudinal cohort of 158 patients with traumatic brain injury and 27 orthopaedic controls, 54% of the brain-injured group developed a mood disorder at some point during the first year, compared with 22% of patients with multiple traumatic injuries without

central nervous system involvement. The comparator is important: each group experienced major trauma, but only the brain-injured group had central nervous system involvement.

9.3 Depression: clinical meaning and anatomical model

Depression after traumatic brain injury is not merely sadness about disability, although loss and demoralisation may contribute. Diagnose a depressive syndrome by integrating mood, loss of interest, biological change, negative cognition, risk, function and course, while interpreting cognitive and somatic complaints in the context of brain injury. Fatigue, slowed thinking, sleep disturbance and reduced activity may arise from several post-injury processes. Their presence should prompt careful phenomenology rather than automatic inclusion or exclusion.

The syndrome has consequences beyond emotional suffering. Once present, post-traumatic depression predicts poorer quality of life, interferes with rehabilitation and adherence, reduces the chance of complete recovery, and increases the risk of comorbid post-traumatic stress disorder and aggression. Thus mood assessment belongs inside the rehabilitation formulation. Reduced participation may be a depressive consequence, a cognitive limitation, apathy, fear-based avoidance or a mixture. The behavioural observation is the beginning of analysis, not the diagnosis.

Left lateral prefrontal involvement provides a useful anatomical correlate without becoming a lesion rule. In one series, post-TBI major depression was associated with reduced grey-matter volume in the lateral aspect of the left prefrontal cortex. A proposed systems account links depression with deactivation of more lateral and dorsal frontal cortex and increased activation of ventral limbic and paralimbic structures. This model explains how disturbed regulation between frontal control systems and affect-generating networks could contribute to depressive experience.

Anatomy must be used proportionately. A left prefrontal finding can support biological plausibility, but absence of that finding does not exclude depression and its presence does not diagnose it. The clinical syndrome remains primary. Professional sport offers another association rather than an individual diagnostic test: professional football players who sustain concussions are 1.5 to 3 times more likely to develop depression than teammates without head injury.

> **GOOD TO REMEMBER**
>
> A patient who disengages from rehabilitation needs a mood formulation, not a character judgement. Ask whether the dominant process is depressive distress, fear-based avoidance, cognitive overload, reduced initiation or treatment toxicity before attributing “poor motivation.”

9.4 Anxiety disorders after traumatic brain injury

Post-TBI anxiety disorders form a spectrum rather than a single undifferentiated complaint. Generalised anxiety disorder may be the most common anxiety disorder after traumatic brain injury. A meta-analysis of 22 studies including 1,145 participants found estimates from 2% to 28%, with a pooled prevalence of **11%**. The range again warns against treating a single figure as universal across settings and injury severities.

Other anxiety presentations show different patterns. Panic disorder has been reported in 6% to 14%, with traumatic brain injury doubling risk; agoraphobia occurred in 7%, also at twice the expected risk. Social anxiety disorder occurred in 5%, twice the rate after other injuries. By contrast, specific phobias were reported at 1% to 10%, similar to the general population, while obsessive-compulsive disorder at 1% to 5% was not significantly elevated. The examination point is discrimination: not every anxiety diagnosis is increased to the same extent.

PRESENTATION	SUPPORTED EPIDEMIOLOGICAL SIGNAL	CLINICAL QUESTION	INTERPRETIVE CAUTION
Depressive disorder	Very wide reported frequency, with greatest risk early and possible persistence	Is there a depressive syndrome rather than isolated fatigue or disengagement?	Somatic and cognitive symptoms need contextual interpretation
Generalised anxiety	Most common anxiety candidate, with a pooled estimate from mixed-severity studies	Is worry pervasive across domains and difficult to disengage from?	Do not collapse all post-injury fear into generalised anxiety
Post-traumatic stress	Higher likelihood after TBI than after other traumatic injuries	Does the presentation organise around traumatic threat and its reminders?	Amnesia for the event is not reliably protective
Psychotic disorder	Reported frequency and timing vary greatly across samples	Are hallucinations or delusions prominent and attributable to injury?	Exclude symptoms occurring only during delirium

Assessment should specify the feared outcome, triggers, mental and bodily response, avoidance, safety behaviours and functional cost. Anxiety may focus on recurrence of injury, bodily sensations, cognitive failure, public embarrassment or traumatic reminders. These themes can coexist, but naming the dominant fear system leads to a clearer diagnosis than recording “anxious after TBI.” Drug choices and cautions belong to the dedicated section on pharmacological management in traumatic brain injury psychiatry; psychological care should follow current syndrome-specific guidance.

9.5 Post-traumatic stress disorder and the amnesia trap

Post-traumatic stress disorder after TBI is neither paradoxical nor rare enough to ignore. Veterans with traumatic brain injury show a threefold greater rate of PTSD than veterans without traumatic brain injury, while civilians with traumatic brain injury have almost twice the likelihood seen after other traumatic injuries. Reported comorbidity spans 12% to 30% in civilians and 12% to 89% in veterans. Population, exposure and ascertainment therefore need to accompany any quoted range.

The classic error is to assume that impaired explicit recall blocks trauma-related illness. Earlier accounts proposed that amnesia for the event would protect against PTSD, but later studies did not support that conclusion. **Implicit trauma memory** offers the bridge: even after moderate or

severe injury, aspects of the traumatic experience may be encoded implicitly and may continue to shape emotion and behaviour. A patient need not provide a coherent autobiographical narrative before trauma cues can provoke distress, arousal or avoidance.

TRAP **Do not write that post-traumatic amnesia protects against PTSD.** Explicit recollection may be incomplete while implicit trauma memory still influences emotional and behavioural responses; moderate or severe TBI does not make PTSD impossible.

Clinically, explore what is remembered, what is reconstructed from others, what remains fragmentary, and what sensory or situational cues trigger a response. Avoid leading the patient toward a narrative that is not genuinely recalled. At the same time, do not demand a detailed narrative as proof of trauma-related pathology. The formulation may rest on cue-linked distress, avoidance, altered threat perception and behavioural change even when episodic recall is limited.

The combination of traumatic brain injury and PTSD appears to produce worse outcomes than either condition alone. This makes dual recognition important: cognitive or behavioural problems should not be attributed exclusively to brain injury when a trauma syndrome is active, and trauma-related symptoms should not obscure neurological limitations. Their combination can shape rehabilitation engagement, family interpretation and functional recovery.

9.6 Psychosis after traumatic brain injury

Psychotic disorder due to TBI is characterised by prominent hallucinations or delusions judged to be a direct physiological consequence of traumatic brain injury, with symptoms occurring exclusively during delirium excluded. This definition demands more than discovering that injury occurred before psychosis. Establish the psychotic phenomena, their chronology, the neurological context, premorbid and family history, substance exposure, current medical state and whether fluctuating confusion better explains the experience.

Reported occurrence varies dramatically by study design. Retrospective review of World War II records found a rate of 7%, whereas a prospective inpatient rehabilitation study found hallucinations and/or delusions at some point after injury in 64% of people with moderate or severe TBI. These are not interchangeable populations or methods. The broad spread should provoke methodological caution, not a midpoint estimate.

Phenomenology is more consistent than the headline rates. Among 71 published schizophrenia-like cases, almost all patients had auditory hallucinations, one in five also had visual hallucinations, 85% had delusions, most often persecutory, and 43% had negative symptoms. The predominance of delusions, paranoia and auditory hallucinations is also described in an alternate clinical account. Visual hallucinations can occur, but their presence should widen assessment rather than automatically disprove the post-traumatic formulation.

Latency is central. The mean onset in the schizophrenia-like case series was **3.3 years after injury**. Another source describes a bimodal distribution, with onset within the first 2 years or after 5 years. A mean can coexist with a bimodal distribution, so the accounts are not formally

contradictory. Together they warn that psychosis may not begin in the immediate recovery period.

Risk evidence further supports vigilance without determinism. A meta-analysis of nine studies involving 180,859 people found subsequent schizophrenia risk elevated after traumatic brain injury, with relative risk 1.65 and a 95% confidence interval from 1.17 to 2.32. Family history of schizophrenia confers particularly heightened risk. An alternate account adds family history of psychosis and prolonged unconsciousness as risk factors, reports equal occurrence after mild and more severe injury, notes that most cases were male, and estimates psychosis in about 20% of severe TBI survivors.

PITFALL

Do not attribute hallucinations to a chronic post-traumatic psychosis while attention, awareness and cognition are fluctuating. Psychotic symptoms occurring only during delirium belong to the delirium formulation; its diagnostic criteria and subtypes are taught in the dedicated section on delirium.

9.7 Secondary mania and the seizure interface

Secondary mania after TBI is less frequent than post-traumatic depression. Estimated bipolar-type disorder rates during the first postinjury year range from 2% to 9%. The syndrome should be distinguished from isolated irritability, impulsivity, sleep disruption, emotional lability or frontal disinhibition. An acquired mood syndrome requires an organising change in mood, with associated activation interpreted as part of that syndrome rather than counted from a behavioural checklist.

In one cohort, manic episodes lasted about 2 months, while elevated, expansive or irritable mood persisted for a mean of 5.7 months. Manic syndromes were significantly associated with hyperactivity, disinhibition, aggression and hypersexuality, but were not related to injury type or severity or to personal or family psychiatric history. The apparently longer mood disturbance than episode duration illustrates why the clinician should describe the syndromal episode as well as the residual affective state.

Secondary mania is associated with post-traumatic seizures, particularly focal-onset seizures with impaired awareness and temporal-lobe epilepsy. Post-traumatic psychosis is likewise associated with seizure disorder in approximately one-third of cases. These associations justify integrating seizure history into psychiatric assessment, but they do not transfer ownership of seizure classification or anticonvulsant choice to this fragment. The psychiatric implications of post-injury seizures are covered in the dedicated section on raised intracranial pressure, hydrocephalus and post-injury seizures; classification and pharmacology belong in the Epilepsy and psychiatry notes.

9.8 Integrated assessment, communication and care boundary

A useful formulation ends in a map of what must be observed, treated and reviewed. The interview should move from chronology to syndrome, from syndrome to attribution, and from

attribution to functional consequences. The following sequence prevents the assessment from becoming a list of post-injury complaints:

- **Baseline:** document premorbid psychiatric illness, personality, function and relevant family history.
- **Course:** reconstruct injury, early neurobehavioural change, symptom onset, remissions and delayed emergence.
- **Phenomenology:** separate depressive, anxiety, trauma-related, manic and psychotic experiences instead of using "behaviour change" as a diagnosis.
- **Mimics and modifiers:** consider fluctuating cognitive states, substances, treatment effects, pain, sleep disturbance and the psychological impact of disability.
- **Function and risk:** assess rehabilitation participation, self-care, relationships, work, self-harm, aggression, vulnerability and the family's capacity to support care.
- **Review:** define what change would confirm improvement, what deterioration requires reassessment, and who will provide longitudinal observation.

MNEMONIC

MAPS

Mood syndromes, including depression and secondary mania; **A**nxiety syndromes; **P**ost-traumatic stress; **S**chizophrenia-like psychosis. MAPS is a recall spine for psychiatric sequelae, not a diagnostic instrument.

Communicate uncertainty honestly. "After TBI" describes chronology; "due to TBI" is a causal formulation that requires support. Explain to families that brain injury can alter vulnerability and expression without making every later emotion neurological or every harmful act inevitable. Repeated review is often more accurate than premature certainty, especially when syndromes overlap or emerge late.

Keep the care boundary explicit. Personality change, aggression and disinhibition are taught in their dedicated section. Cognitive impairment, post-traumatic amnesia, rehabilitation and chronic traumatic encephalopathy are taught separately. Pseudobulbar affect is cross-referred to the dedicated section on its phenomenology and treatment. Drug choices, doses, cautions and seizure-threshold considerations belong to the dedicated section on pharmacological management in traumatic brain injury psychiatry, with general antidepressant pharmacology in the Mood disorders: pharmacotherapy notes.

After traumatic brain injury, diagnose the syndrome, reconstruct its chronology, test causal attribution, and keep depression, anxiety, PTSD, psychosis and secondary mania distinct even when they coexist.

10 · Personality change, aggression and disinhibition after traumatic brain injury

After traumatic brain injury, a changed person may be the presenting complaint even when the family cannot name a psychiatric syndrome. The clinically important change may involve social judgement, impulse control, emotional expression, motivation or aggressive behaviour. The psychiatrist must convert a moralised account such as "he has become rude" into an operational description of what changed, when it changed, what triggers it and which control function is failing.

IN INDIA

Lead with behavioural management, family psychoeducation and environmental structure; where medication is used for aggression, choose a familiar agent at the lowest effective dose and review the need regularly.

This topic earns marks through distinctions. Personality change is broader than aggression; disinhibition is a failure of restraint; agitation is not synonymous with aggression; and apathy is a low-output disorder that may be mistaken for depression. A strong answer links phenomenology to prefrontal circuitry, uses collateral history to establish change from baseline, assesses risk without equating all irritability with violence, and keeps pharmacological detail within the dedicated section on pharmacological management in traumatic brain injury psychiatry.

COMPORTMENT	DISINHIBITION	AGGRESSION	APATHY
SOCIAL DECISION-MAKING	FAILED RESTRAINT	THREAT OR INJURY	FAILED INITIATION

10.1 The core construct: acquired change in personality and comportment

Post-traumatic personality change is an acquired alteration in the patient's characteristic manner of relating, deciding and regulating behaviour after brain injury. It is best established as a difference from the premorbid person, not as a free-standing adjective applied during a difficult interview. The family may describe a newly abrupt, suspicious or argumentative manner. These descriptions become clinically useful only after the examiner asks for concrete episodes, context, consequences and comparison with the previous baseline.

Comportment refers to the regulation of behaviour in its social and emotional setting. The relevant prefrontal circuitry modulates decisions according to reward and punishment, helps generate and express emotion, and supports insight into one's own conduct. Injury can therefore leave basic language or motor capacity available while degrading the ability to choose what is appropriate, anticipate interpersonal consequences or recognise that one's conduct has changed.

This explains the apparent paradox often seen in consultation. The patient may understand the words of a rule, recall that relatives objected, and still repeat the behaviour when an immediate cue or frustration appears. Knowledge of the rule and online use of the rule are different capacities. The examiner should therefore document whether the failure concerns understanding, memory, initiation, inhibition, emotional modulation or insight rather than compressing all of them into "personality disorder."

RULE **Describe acquired behaviour in operational terms before assigning motive or diagnosis.** Establish the premorbid baseline, identify the failed control function, and show how the change affects safety, relationships and ordinary roles.

The word acquired matters. A longstanding interpersonal style may shape the post-injury presentation, but it is not itself evidence of injury-induced change. Conversely, a dramatic departure from baseline may be missed if the clinician relies only on the patient, because impaired insight is part of the relevant circuitry. The formulation must hold premorbid traits and new neurological change together without allowing either to erase the other.

10.2 Neuroanatomical logic of disinhibition

Kaufman's Clinical Neurology for Psychiatrists links personality change after traumatic brain injury with frontal and temporal damage. When frontal inhibitory systems are impaired, behaviour that would ordinarily be checked can reach expression. The resulting **disinhibition** may appear as unbridled aggressiveness, excessive talk, impulsivity, hyperactivity or a more generally unrestrained interpersonal manner.

Disinhibition is therefore not simply "too much behaviour." It is behaviour insufficiently filtered by context, future consequence and social meaning. A patient may speak at length despite cues to stop, act on an immediate urge before considering risk, or escalate rapidly after a minor frustration. The shared mechanism is weak restraint, although the surface behaviours differ. This functional description is more useful than assuming that every activated presentation is a primary mood episode or that every offensive act is deliberate.

The anatomical bridge is sharpened by the well-known case of **Phineas Gage**. Modern imaging of the injury demonstrated damage involving orbitofrontal and medial prefrontal regions, areas contributing to reward-punishment decisions, emotional processing and insight. The case is memorable because it illustrates how disruption of control circuitry can alter social conduct without reducing the person to a simple motor or language deficit.

Localisation should still remain proportionate. A behaviour is not a scan, and a colourful symptom does not establish a precise lesion. Frontal and temporal injury may be distributed, and several behavioural functions may fail together. Use the anatomical account to explain the pattern, not to claim a focal site from one impulsive incident. The dorsolateral, orbitofrontal and medial frontal syndromes in full belong in the dedicated section on frontal lobe syndromes and their subtypes.

MNEMONIC

BRAKE

Baseline personality; **R**eward and punishment decisions; **A**ffect generation and expression; **K**eeping impulses in check; **E**valuating one's own conduct through insight. BRAKE is a clinical reading frame for comportment, not a diagnostic instrument.

10.3 Phenomenology: separate the behavioural syndromes

The phrase "behavioural disturbance" is too coarse for either an examination answer or a clinical formulation. The same ward report may contain pacing, shouting, threats, impulsive touching, repetitive demands, emotional indifference and failure to begin self-care. These observations do not all name the same construct. The table separates them by the dominant failure and by what must be demonstrated before using the label.

CONSTRUCT	CORE CLINICAL MEANING	TYPICAL DESCRIPTION AFTER TBI	QUESTION THAT SEPARATES IT
Personality change	Acquired alteration in characteristic manner, potentially abrupt, suspicious or argumentative	A sustained departure from premorbid interpersonal style	What is genuinely new compared with baseline?
Comportment change	Impaired modulation of reward-punishment decisions, emotional expression and insight	Poor social judgement despite apparently intact basic abilities	Can the patient use context and consequence to guide conduct?
Disinhibition	Loss of inhibitory control with aggressiveness, loquaciousness, impulsivity or hyperactivity	Behaviour escapes restraint when an urge or cue appears	Did the person understand the boundary but fail to inhibit the response?
Agitation	Excessive, usually nonproductive motor activity associated with inner tension	Restless motor output that need not threaten anyone	Is there excessive motor activity, threat, injury, or a combination?
Aggression	Overtly threatening or injurious conduct toward a person or object	Usually reactive and impulsive in the setting of TBI-related cognitive impairment	What was threatened or harmed, and what immediately preceded it?
Apathy	Reduced motivation, initiation, interest, activity, emotional responsiveness or spontaneity	Low self-generated behaviour, sometimes without subjective distress	Is the main failure distress, desire, initiation, or motor capacity?

Several constructs may coexist. A patient can be apathetic for much of the day yet become impulsively aggressive when pressed to perform a task. Another may pace from inner tension without threatening anyone. A third may make offensive comments because inhibitory control is poor but never become physically aggressive. The formulation should preserve these distinctions rather than choosing the most dramatic label as the whole diagnosis.

Descriptions should also separate frequency from severity and intent from effect. A brief threat can be clinically serious; repeated restless movement may be disruptive without being aggressive; an impulsive act can cause major injury without having been planned. Neutral behavioural

language protects accuracy and dignity. It also gives the team something observable to track rather than inviting arguments about whether the patient is "really difficult."

10.4 Agitation is not aggression

Agitation is excessive, usually nonproductive motor activity accompanied by inner tension. Its centre of gravity is motor overactivity. **Aggression** is overtly threatening or injurious behaviour directed toward another person or an object. Its centre of gravity is threat or harm. The patient may have either construct alone or the constructs may occur together.

This distinction changes observation. For agitation, record what the body is doing, whether the movement serves a purpose, the apparent tension and the environmental setting. For aggression, document the target, the threatening or injurious act, the immediate antecedent, the consequence and how the episode ended. "Agitated and aggressive" without this detail hides the very information needed for formulation and prevention.

In traumatic brain injury with cognitive impairment, aggressive behaviour tends to be **reactive and impulsive aggression** rather than proactive aggression. The episode commonly reads as a rapid response to frustration, demand, misunderstanding or immediate provocation rather than as planned instrumental harm. This does not make the behaviour harmless or remove responsibility for safety. It changes the causal question from "What long-term goal did the aggression serve?" to "What antecedent met an impaired regulatory system?"

TRAP **Do not use agitation and aggression as interchangeable words.** Pacing with tension is not automatically violence, while a threat or injurious act may occur without sustained motor overactivity. Define the observed construct before discussing cause.

Overt Aggression Scale may be used to track response. Its value here is conceptual as well as practical: repeated observation should be organised around overt behaviour rather than global impressions. Detailed instrument administration and scoring belong to the Psychotherapies, clinical assessment, MSE and nosology notes. In this section, the essential point is to use consistent behavioural anchors across observers and across time.

10.5 Why aggression occurs: a layered risk formulation

Post-traumatic aggression does not arise from lesion location alone. The supported risk formulation combines premorbid factors, the location of injury and the sequelae produced by the injury. Premorbid poor social functioning is relevant, and substance misuse has been associated in some studies. Frontal-lobe injury carries greater risk than injury to other brain regions. The most powerful risk factor in the clinical account is **TBI-induced depression**.

The depression association is not a decorative comorbidity. During the first year after traumatic brain injury, significant aggressive behaviour was found in **17 of 30 patients with major depression, or 57%**, compared with **10 of 44 patients without a mood disorder; p = .004**. The comparison explains why mood assessment belongs inside an aggression formulation rather than after it.

The risk factors answer different questions. Premorbid functioning describes the regulatory and social context brought into the injury. Lesion location indicates which control systems may have been disrupted. Post-injury sequelae describe new states that can amplify irritability, distress or behavioural dyscontrol. A useful formulation shows how these layers converge in this patient instead of presenting them as a memorised list.

- **Baseline:** establish premorbid social functioning, usual conflict style and any substance-related context without assuming that history determines the present episode.
- **Brain injury:** identify whether frontal involvement and impaired inhibition plausibly fit the observed reactive pattern.
- **Current sequelae:** assess depression and other post-injury problems that may alter tolerance, judgement or emotional regulation.
- **Antecedents:** reconstruct what happened immediately before the behaviour, including demands, frustration, misunderstanding and environmental overstimulation.
- **Consequences:** document harm, near-harm, family burden, occupational disruption and what responses appear to escalate or settle the episode.

The risk formulation must remain dynamic. A historical vulnerability is not a permanent prediction, and an absence of previous violence is not proof of current safety. Likewise, frontal injury is a risk factor, not a behavioural destiny. The clinician's job is to integrate enduring background with present state, immediate triggers and observed control, then revise the formulation when the pattern changes.

GOOD TO REMEMBER

When aggression appears after traumatic brain injury, assess mood early. Depression may be clinically less dramatic than the behavioural episode yet more important to the risk formulation than a long catalogue of remote vulnerabilities.

10.6 Apathy: the low-output counterpoint

Apathy after traumatic brain injury involves impairment of motivation, initiation, interest, activity, emotional responsiveness or spontaneity. The patient may not begin tasks, may contribute little without prompting, or may show reduced emotional engagement. The central observation is diminished self-generated behaviour. It should not be confused with inability to understand a task or inability to perform the required movement.

Apathy is especially easy to mislabel as depression. Patients with apathy are often not distressed by their reduced activity, whereas the clinician may infer sadness simply from sparse speech, inactivity or blunted responsiveness. Depression and apathy can coexist, but they are not synonyms. Ask whether the patient wants to act, feels distressed by the loss, initiates when structure is removed, responds to prompting and shows broader depressive experience. The distinction matters because an apparent low-output syndrome and an affectively painful syndrome require different formulations.

Imaging evidence links post-traumatic apathy with damage in the left hemisphere, particularly the frontal lobe, anterior cingulate, insula, corpus callosum and corona radiata. The distributed list fits the clinical idea that motivation and initiation depend on connected systems rather than a single isolated point. Laterality must be retained accurately: the reported association is left-sided.

Apathy also clarifies why "personality change" should not be treated as synonymous with disinhibition. Traumatic brain injury can produce a high-output failure of restraint, a low-output failure of initiation, or a mixed picture. In a mixed presentation, behaviour may be sparse until an external demand appears, after which the patient reacts impulsively. That is not inconsistency. It indicates that initiation and inhibition are separable control functions.

PITFALL

Do not diagnose depression solely because the patient is inactive, emotionally unresponsive or dependent on prompts. Establish subjective distress and the broader mood syndrome, while separately describing motivation, initiation and spontaneity.

10.7 Bedside assessment and documentation

The decisive evidence is usually a triangulation of patient interview, direct observation and collateral history. Begin with chronology: the premorbid interpersonal style, the injury, the earliest observed change, the subsequent course and the current pattern. Ask relatives for episodes rather than labels. "Argumentative" should become a description of what was said, what preceded it, whether correction altered the behaviour and whether the patient later understood its effect.

Direct observation should be equally operational. Note initiation of conversation, ability to wait, response to limits, modulation of voice, interruption, persistence after feedback and recognition of interpersonal cues. An apparently calm structured interview does not exclude dyscontrol in a noisy home or demanding ward. Conversely, conflict during one interview does not by itself establish acquired personality change. Context is part of the phenomenon.

Risk assessment should distinguish current intent from impaired control. Ask directly about recent threats or injury, targets, access to means, substance use around episodes, escalating frequency, warning signs and the family's ability to maintain safety. Then examine the reactive sequence: antecedent, interpretation, impulse, act and consequence. This sequence identifies where prevention may be possible without pretending that a behavioural formulation can guarantee prediction.

Documentation should use neutral verbs. Record "paced continuously," "struck the door after a limit was set," or "continued speaking despite repeated cues" rather than "was impossible" or "behaved badly." State what came from collateral history and what was observed directly. Describe insight as specific to the event: the patient may recognise that an act occurred yet fail to appreciate why it was dangerous or how it affected others.

Track patterns consistently. Where aggression is the target, the Overt Aggression Scale can support repeated measurement. Use the score alongside the narrative of antecedents, targets and

consequences. Broader cognitive testing may clarify associated impairment, yet the instruments and their scoring are taught in the Psychotherapies, clinical assessment, MSE and nosology notes. The behavioural formulation remains grounded in change from baseline and conduct in context.

10.8 Synthesis, communication and treatment boundary

A concise diagnostic synthesis should state the acquired behavioural change, dominant control failure, anatomical coherence, relevant risk factors and degree of insight. For example, the formulation may describe a new reactive pattern with poor inhibition after frontal injury, amplified by a current depressive syndrome, while acknowledging premorbid social vulnerability. This says more than "aggression due to TBI" because it identifies the mechanism that needs to be tracked.

Communication with relatives should validate paired realities: the behaviour may arise from injured control systems, and its consequences still require firm safety boundaries. Neurological explanation is not moral condemnation, but neither is it permission to ignore harm. Families often need help separating the person from the behaviour while continuing to record antecedents, recognise escalation and seek urgent help when safety cannot be maintained.

Do not let anatomy become fatalism. Risk factors describe probability and mechanism, not inevitability. The patient retains capacities that can be supported, environments can change the load placed on inhibitory systems, and mood or other sequelae may be clinically modifiable. The present fragment establishes the target syndrome and the reasoning needed to distinguish its components. Detailed medication selection, doses, cautions and the lowered seizure-threshold problem belong to the dedicated section on pharmacological management in traumatic brain injury psychiatry.

For an examination answer, end where the evidence ends. Define personality and comportment change, explain disinhibition through frontal control failure, distinguish agitation from aggression, characterise reactive impulsive aggression, present the layered risk formulation with depression prominent, and contrast apathy with depression. Then add assessment through collateral history, operational observation and repeated behavioural tracking. This sequence is complete without borrowing treatment details or reproducing the full frontal-lobe syndrome classification.

> After traumatic brain injury, describe whether the dominant failure is social judgement, inhibition, threat control or initiation; do not collapse personality change, disinhibition, agitation, aggression and apathy into one label.

11 · Cognitive impairment, post-traumatic amnesia, rehabilitation and chronic traumatic encephalopathy

The marks lie in following cognition across time, not in treating every post-injury memory complaint as the same disorder. **Post-traumatic amnesia** is an evolving recovery period in which new learning is prominently impaired while memory for the trauma and the immediately

surrounding past is also disturbed. It must be distinguished from the confusional state that can precede it and the executive retrieval problem that can follow it, and from the much later question of neurodegeneration after repeated injury.

This distinction has practical force. The duration of amnesia carries unusual prognostic weight, rehabilitation targets change as the dominant cognitive problem changes, and apparent recovery of conversation does not prove restoration of memory or judgement. At the chronic end, a history of traumatic brain injury can increase later cognitive risk without establishing chronic traumatic encephalopathy. A strong answer therefore moves from definition and bedside interpretation to prognosis, rehabilitation, later dementia risk and the strict pathological boundary around CTE.

Strongest	Within 6 months	Until 18 months	Autopsy only
PTA DURATION AS OUTCOME PREDICTOR	MOTOR AND LANGUAGE MAXIMUM	INTELLECTUAL RECOVERY MAY NOT PEAK	DEFINITIVE CTE DIAGNOSIS

11.1 What post-traumatic amnesia means

PTA begins to make clinical sense when arousal and simple attention have relatively recovered, yet the patient remains unable to lay down a stable record of ongoing experience. Its core is impaired formation of new memories after injury, accompanied by loss of memory immediately before or during the event. The apparent alertness can deceive: a patient may answer, converse and engage briefly while later retaining little of what occurred. This is why ordinary social responsiveness cannot be used as proof that amnesia has ended.

The mnemonic disturbance reflects injury to hippocampal systems that support memory function. This is a functional explanation, not permission to infer a focal structural lesion in every patient. The clinician should identify whether information is being consistently retained across encounters, whether orientation is becoming stable, and whether the account of the event is contracting or changing during recovery. Collateral observation is particularly valuable because the patient cannot reliably report a period that was not successfully encoded.

The **Galveston Orientation and Amnesia Test** can assess the duration and severity of PTA. It belongs here as the named instrument for this construct, but its item content, administration and scoring should be learned from the authorised instrument source and are not reproduced here. General teaching on standardised cognitive instruments belongs in the Psychotherapies, clinical assessment, MSE and nosology notes.

RULE **Do not equate wakefulness with recovery from post-traumatic amnesia.** Relative return of arousal and simple attention may coexist with a continuing inability to form new memories, so repeated observation of retention matters more than conversational fluency.

11.2 Retrograde loss, anterograde failure and the bedside question

Retrograde amnesia covers memory loss for the trauma and immediately preceding events. This retrograde gap tends to diminish during recovery, so the patient may progressively recover portions of the lead-up to the injury. The clinical question is therefore not simply whether a gap

exists, but whether its boundary is changing and whether the reconstructed account comes from genuine recall, collateral information or later inference.

Anterograde amnesia concerns newly presented information. It is uncommon in mild traumatic brain injury but very common after moderate-to-severe traumatic brain injury. The essential mechanism is failed encoding: information never entered a durable memory trace. Consequently, the anterograde period itself is never subsequently recovered. A patient may later learn what others report happened, but this acquired narrative is not recovery of the original experience.

A compact assessment keeps the following questions separate:

- **Past boundary:** What is missing from the trauma and the immediately preceding period, and is that boundary shrinking?
- **New retention:** Does information discussed now remain available after a delay and across subsequent encounters?
- **Source:** Is the later account recalled, learned from others or reconstructed from context?

This separation prevents a common interpretive error. Improvement in retrograde recall does not prove recovery of new learning, and an orderly narrative supplied after the event does not show that the person encoded the original period. Assessment is longitudinal because PTA is a recovery state, not a single interview finding.

11.3 The cognitive recovery sequence

Acute traumatic encephalopathy can be understood as a changing hierarchy. After emergence from coma, a persistent vegetative state or a minimally conscious state, the patient may enter a **post-traumatic confusional state**. This is delirium-like, with acute and fluctuating impairment of attention, perceptual disturbance, delusional thought, psychomotor agitation or retardation, and affective lability. The full delirium syndrome and its diagnostic criteria are taught in the dedicated delirium section; here the point is its position in post-injury recovery.

RECOVERY PHASE	DOMINANT DISTURBANCE	MEMORY INTERPRETATION	CLINICAL IMPLICATION
Disorder of consciousness	Impaired level of consciousness	Memory cannot be interpreted independently of arousal	Establish the level and trajectory of consciousness first
Post-traumatic confusional state	Fluctuating attention, perception, thought, psychomotor activity and affect	Inconsistent registration occurs within a broader confusional state	Do not reduce the presentation to isolated amnesia
Post-traumatic amnesia	Failure to form new memories despite relative recovery of arousal and simple attention	Encoding failure dominates	Track retention and the changing amnestic boundary
Post-PTA dysexecutive phase	Executive dysfunction associated with frontal-subcortical injury	Retrieval failure becomes more important than encoding failure	Shift rehabilitation toward executive control and compensation

After PTA resolves, **post-traumatic dysexecutive syndrome** tends to predominate because closed head injury has a predilection for frontal-subcortical circuitry. The patient may still describe "poor memory," but the mechanism has changed. During PTA, the central problem is failure to encode; later, information may have been stored but cannot be efficiently organised, searched for or retrieved. This distinction explains why a single complaint can require different assessment and rehabilitation strategies at different points in recovery.

GOOD TO REMEMBER

Ask what kind of memory failure is occurring now. Repetition without later retention suggests continuing encoding failure; knowledge that is inconsistently accessed after PTA points toward executive retrieval difficulty. The label "memory problem" is too broad to plan rehabilitation.

11.4 Cognitive pattern, localisation and awareness

The location of trauma correlates inconsistently with the resulting cognitive impairment. Traumatic injury often does not respect the tidy localisation expected from a single focal lesion, so a scan location should not be made to explain every deficit. The important grounded exception is left temporal lobe injury, which routinely produces vocabulary difficulty resembling anomic aphasia. This is a useful localising clue, not a template for converting all post-TBI cognitive findings into focal lobe syndromes.

Awareness is equally important. Patients after TBI, especially those with right hemisphere involvement, may show **anosognosia** or denial of their deficits. The problem is not mere disagreement with staff. Unawareness obstructs realistic goal setting, interferes with rehabilitation and may expose the patient to further traumatic brain injury. A person who does not recognise impaired judgement or memory cannot reliably calibrate risk, assistance or readiness to resume demanding roles.

Assessment should therefore triangulate the patient's account, observed performance and collateral information. Discrepancy is clinically informative, but it must not be used as a moral verdict. Some patients underestimate deficits because awareness systems are impaired; others accurately describe failures not obvious in a short interview. The psychiatrist's task is to define the discrepancy, test its functional consequences and build goals that preserve dignity while acknowledging risk.

PITFALL

Do not interpret confident self-report as evidence of intact capacity for rehabilitation decisions. Post-TBI unawareness can be part of the cognitive syndrome, particularly with right hemisphere involvement, and may make the most certain patient the least reliable judge of current limitations.

11.5 Prognosis and the shape of recovery

The duration of PTA is a more reliable predictor of overall neuropsychological outcome than either the depth or duration of coma. It predicts the broad cognitive and functional burden because it samples the period during which organised new learning remains disrupted. This does

not turn duration into destiny. It gives the clinician a stronger prognostic anchor than the dramatic appearance of the earliest coma alone.

Recovery is domain-specific rather than synchronous. Motor and language abilities usually reach their maximum earlier, whereas intellectual recovery may peak substantially later, as summarised in the statistics strip. This lag matters when families assume that improved walking or fluent speech means cognition has also plateaued. It also argues against declaring the final intellectual outcome too early merely because visible neurological recovery has slowed. Older patients generally recover more slowly and less completely than younger patients.

Prognosis is modified by more than intracranial pathology. Adverse associations include older age, lower premorbid cognitive reserve, poorer socioeconomic status, limited social support, involvement in legal or work-compensation proceedings, and neurological or psychiatric conditions arising before or after injury. The APOE-epsilon4 allele has also been associated with poorer clinical outcome. These are modifiers, not automatic explanations for an individual's disability and not grounds for dismissing symptoms.

The practical formulation should state the injury-related anchor, the present cognitive phase, domain-specific strengths and weaknesses, awareness, psychiatric comorbidity, social supports and real-world participation. Prognosis is then expressed as a trajectory with uncertainty. This is more defensible than a fixed promise and more useful than the vague statement that “recovery varies.”

TRAP **Do not use coma duration as the best cognitive prognostic marker when PTA duration is available.** Coma is visually dramatic, but the duration of post-traumatic amnesia is the more reliable predictor of overall neuropsychological outcome.

11.6 Rehabilitation: locus, focus and modus

Cognitive rehabilitation after traumatic brain injury is primarily multidisciplinary and nonpharmacological. The overall purpose is not simply to improve a test score. It is to restore intellectual functioning where possible, strengthen remaining abilities, build compensatory mechanisms and increase safe participation. Treatment must follow the changing cognitive phase: an approach suited to confusion or failed encoding will not be sufficient when the dominant barrier is later executive retrieval.

LOCUS. Early care belongs in a setting that can manage impaired consciousness, confusion, agitation and safety while repeatedly reassessing recovery. As medical stability and participation improve, rehabilitation can move through structured inpatient or outpatient work into the patient's ordinary environment. Locus is chosen by supervision needs, functional dependence, behavioural stability, caregiver capacity and access to the relevant disciplines, not by diagnosis alone.

FOCUS. During the confusional and amnestic phases, the immediate targets are safe engagement and recognition of what information is not being retained. After PTA, focus shifts toward executive organisation, retrieval, judgement, awareness and the barriers that prevent therapy

participation. Across recovery, depression, anxiety and insomnia require active identification and treatment because they can worsen distress, efficiency and engagement. Long-term goals should be framed in observable function and participation rather than in the disappearance of every subjective complaint.

MODUS. The grounded rehabilitation package includes:

- **Restorative disciplines:** cognitive, behavioural, physical, occupational and speech rehabilitation, coordinated around shared functional goals.
- **Psychiatric and sleep care:** identification and treatment of depression, anxiety and insomnia, integrated with rather than separated from rehabilitation.
- **Social and compensatory work:** peer interaction, enhancement of remaining functions and development of compensatory mechanisms for deficits that persist.

Medication is adjunctive and the evidence base is limited. Methylphenidate and other dopamine-enhancing medicines may improve attention and thereby directly or indirectly support memory; anticholinesterases may improve memory in patients with severe deficits. These statements support cautious target-based consideration, not a dose, hierarchy or routine prescription. Detailed choice, adverse-effect cautions and monitoring belong in the dedicated section on pharmacological management in traumatic brain injury psychiatry.

Current head-injury and neurorehabilitation guidance should determine transfer criteria, service intensity and local prescribing safeguards. The examination-quality principle is stable: select the setting by need, select the target by phase, and combine cognitive, behavioural, physical, occupational, speech, psychiatric and social modes rather than waiting for a drug to restore cognition.

MNEMONIC

TRACE

Track PTA duration; **R**ecognise retrieval failure after encoding failure; **A**ssess awareness of deficit; **C**oordinate multidisciplinary compensation and restoration; **E**nforce the postmortem boundary for definitive CTE diagnosis.

11.7 Traumatic brain injury and later dementia risk

Long-term decline after TBI must be handled probabilistically. In survivors of moderate TBI who carry two ApoE4 alleles, the risk of Alzheimer disease may be **double**; after severe TBI in the same genetic context, it may be **four times** the risk. Studies of veterans found a **60% increase** in dementia risk among those with a TBI history compared with veterans without such a history.

These associations do not mean that every later cognitive complaint is degenerative, that genotype fixes an individual's outcome, or that dementia after TBI is synonymous with CTE. They mean that previous injury belongs in the risk formulation. A patient with progressive decline still requires the ordinary diagnostic approach rather than a shortcut from history to

label. The dementia syndrome and its full diagnostic workup belong in the Dementias and neurocognitive disorders notes.

Clinically, separate the following questions. What residual deficit has remained relatively stable since the injury? What function is changing now? What alternative neurological, psychiatric, sleep-related, medical or medication-related explanation could account for that change? This temporal analysis protects against opposite errors: overlooking a progressive disorder because "the patient always had TBI," and diagnosing neurodegeneration from a static acquired deficit.

11.8 Chronic traumatic encephalopathy: phenotype, pathology and diagnostic boundary

Chronic traumatic encephalopathy is a neuropathological diagnosis. Definitive diagnosis is currently possible only at autopsy. The defining lesion is irregular phosphorylated tau aggregated in neurons and astrocytes around small vessels at the depths of cortical sulci. Every component matters: phosphorylated tau, neuronal and astrocytic involvement, a perivascular distribution, and concentration at sulcal depths. Generic "tau deposition" is too imprecise for the defining pathology.

The clinical language has changed across contexts. **Dementia pugilistica**, the classic form described in professional boxers, emphasises parkinsonism and nonprogressive cognitive deficits. The modern CTE phenotype places greater emphasis on prominent psychiatric symptoms and progressive cognitive deficit, while parkinsonism can occur in both.

AXIS	CLASSIC ACCOUNT	MODERN ACCOUNT	INTERPRETIVE LIMIT
Historical setting	Professional boxers	Broader progressive neurodegeneration related to repeated TBI	Exposure history supports suspicion but is not diagnostic
Cognitive course	Primarily nonprogressive deficits	Deficits tend to progress	No distinguishing cognitive pattern has been established
Clinical emphasis	Parkinsonism with cognitive deficit	Prominent psychiatric symptoms	Psychiatric symptoms are not a neuropathological diagnosis
Motor syndrome	Parkinsonism occurs	Parkinsonism also occurs	Its presence does not separate the forms

The epidemiological signals explain concern but do not create an in-life diagnostic test. Depending on career length, boxers have an approximate **20% risk of dementia**. Among professional National Football League players, mortality from neurodegenerative disease is **three times higher** than in the general population of the United States. Before death, affected players may develop memory loss, depression, aggressive behaviour and executive dysfunction; autopsy studies describe excess cerebral tau with plaques and tangles.

The diagnostic trap is overreach. No distinctive pattern of cognitive deficit has been validated for CTE, and a history of TBI is not specific because TBI also increases the risk of conventional dementia. Therefore, a living patient with repeated injury, psychiatric symptoms and decline may

warrant careful assessment, longitudinal observation and investigation for established causes, but cannot be given a definitive neuropathological diagnosis from that constellation. Current evidence does not supply an operational in-life clinical pattern that can replace the postmortem standard.

PTA is primarily an encoding failure and the later dysexecutive phase a retrieval problem; definitive CTE remains a postmortem diagnosis.

12 · Pharmacological management in traumatic brain injury psychiatry

After traumatic brain injury, the safest prescription is the one tied to a defined target, chosen for cognitive tolerability, and reviewed often enough to be stopped. Pharmacological management after brain injury is not a list of preferred drugs. It is a method of matching a symptom target to the phase and setting of care while protecting attention, memory, arousal and seizure safety. The examiner therefore expects more than an agent name: state the target syndrome or behaviour, distinguish acute control from chronic treatment, explain the evidence limit, and specify how iatrogenic worsening will be detected.

IN INDIA

SSRIs are a sensible first choice for mood and anxiety after traumatic brain injury; prefer agents that do not lower the seizure threshold and avoid a heavy anticholinergic load, and start any dopaminergic agent for apathy under specialist advice.

Nonpharmacological care should be prioritised, with medication used as an adjunct rather than a substitute. This principle keeps rehabilitation, psychological treatment, environmental modification and family work visible even when the question asks specifically about drugs. The psychiatric syndromes and behavioural phenotypes themselves are described elsewhere in the traumatic brain injury sections; the present task is their TBI-specific pharmacological management.

SSRIs	Beta-blockers	Prodopaminergic	Adjunctive
FIRST-LINE FOR DEPRESSION	MORE RELIABLE FOR AGGRESSION	TRIAL OPTION FOR APATHY	ROLE OF MEDICATION

12.1 The TBI-sensitive prescribing frame

TBI-sensitive prescribing begins by locating the patient. In emergency and inpatient care, the immediate focus is dangerous agitation, severe distress and behaviour that prevents essential assessment or treatment. In rehabilitation, the focus shifts toward participation, irritability, disinhibition, mood and cognitive efficiency. In outpatient care, longer-term depression, anxiety, trauma-related symptoms, psychosis, secondary mania or apathy may become the treatment target. The locus changes what can be observed, how rapidly deterioration can be detected, and which adverse effects are acceptable.

LOCUS	FOCUS	MODUS	PRESCRIBING QUESTION
Emergency setting	Acute agitation, distress or immediate risk	De-escalation, correction of drivers, environmental control; medication only when needed	Will the drug obscure neurological or cognitive reassessment?
Inpatient unit	Stabilisation and safe participation in care	Nursing structure, family input, rehabilitation measures and a target-specific drug trial	Is the apparent response therapeutic benefit or merely sedation?
Rehabilitation setting	Engagement, behavioural control and functional recovery	Cognitive and behavioural rehabilitation with adjunctive medication	Does the medicine improve the target while preserving learning?
Outpatient setting	Persistent affective, anxiety, behavioural or motivational symptoms	Psychotherapy, social and vocational rehabilitation, family work and rational pharmacotherapy	Can benefit be demonstrated in ordinary function?
Refractory illness	Severe symptoms despite adequate care	Reformulation, specialist review and an appropriate procedural option	Has treatment resistance been established without avoidable polypharmacy?

The practical sequence is simple. Define an observable target, record its functional consequence, remove or reduce medicines that may be producing the same problem, and select a single intervention whose risks fit the injured brain. Begin cautiously and increase slowly, but still allow an adequate dose and duration before declaring failure, tapering or adding an adjunct. “Start low and go slow” is not permission to leave the patient indefinitely on an ineffective token dose.

RULE **Treat a target, not the history of head injury.** Name the symptom or behaviour, choose the least cognition-disrupting reasonable option, and judge benefit against function rather than quietness alone.

12.2 Before the prescription: target, burden and safety

The pre-prescribing assessment asks whether medicine is necessary now, whether a current medicine is contributing, and what observation will count as improvement. A target such as “less aggression during personal care” is more useful than “better behaviour.” Baseline arousal, attention, memory efficiency, sleep, motor restlessness and day-to-day function should be described in clinically meaningful terms. This makes later deterioration visible and prevents sedation from being mistaken for recovery.

Brain-injured patients require frequent reassessment for adverse effects and drug interactions, with particular restraint toward polypharmacy and medicines that interfere with attention or memory. Anticonvulsants, antipsychotics and minor tranquillisers may themselves impair cognition; this does not make every use wrong, but it raises the threshold for adding them

without a defined indication. Their general pharmacology belongs in the Epilepsy and psychiatry notes and the Schizophrenia: management, antipsychotics and adverse effects notes.

MNEMONIC

TARGET

Target one symptom or behaviour; **A**ssess baseline cognition and function; **R**educe contributory medicines; **G**ive one rational trial; **E**valuate benefit and harm repeatedly; **T**aper what has no continuing indication.

A defensible monitoring plan therefore includes the following elements:

- **Target:** the syndrome, behaviour or functional barrier being treated.
- **Benefit:** the change expected in distress, participation, safety or ordinary performance.
- **Harm:** cognitive dulling, excessive sedation, neurological adverse effects, behavioural activation or treatment-emergent restlessness.
- **Decision:** continue, adjust, taper or reformulate at a planned review rather than accumulating indefinite adjuncts.

Current institutional and national prescribing guidance should be checked for ordinary contraindications, interaction screening and agent-specific monitoring. No named guideline or TBI-specific monitoring schedule is supplied by the permitted evidence, so neither should be reconstructed from memory.

12.3 Post-TBI depression: first-line treatment and the evidence paradox

Post-TBI depression pharmacotherapy usually begins with an SSRI. Sertraline is the most studied SSRI in this population, yet the major sertraline trials did not establish efficacy. This is the central evidence paradox: clinical first-line status reflects overall safety and tolerability as well as efficacy evidence, while the trial record remains mixed. A good answer preserves this tension and does not convert "most studied" into "proven best."

The prevention evidence illustrates why dose figures require context. A subacute-recovery trial of post-TBI depression prevention used **sertraline 100 mg/day** in 94 patients. The placebo group had four times the hazard of developing depression, and the **number needed to treat was 5.9**. A separate subacute-recovery trial of post-TBI depression prevention using sertraline 50 mg/day for 3 months found no significant difference from placebo. A post-TBI depression treatment vignette also began sertraline at 50 mg/day. These are trial and vignette doses for post-TBI depression, not a recommended target range.

Selection among antidepressants should favour minimal anticholinergic burden, minimal sedation and minimal lowering of the seizure threshold. Tricyclic antidepressants are not preferred after TBI because their anticholinergic effects and seizure-threshold lowering are particularly disadvantageous. Serotonin-norepinephrine reuptake inhibitors can also lower the seizure threshold and should be used at the lowest effective dose. The general comparative

pharmacology and routine selection of antidepressants belong in the Mood disorders: pharmacotherapy notes.

Small randomised studies suggest that methylphenidate may improve depression, cognition and fatigue. That broader possible benefit may be relevant when the clinical target spans affective and cognitive complaints, but small-study signals should not be presented as equivalent to an established first-line antidepressant strategy. The choice still requires a defined target and repeated assessment of function.

TRAP **Do not quote the sertraline trial doses as a TBI therapeutic range.** The evidence supplies doses used in prevention trials and a treatment vignette, but it does not supply a recommended titration schedule or target range for post-TBI depression.

12.4 Refractory depression and neurostimulation

Electroconvulsive therapy after TBI is not contraindicated merely because traumatic brain injury has occurred. It may be considered for refractory post-TBI depression, provided technique is adapted to reduce cognitive and seizure-related burden. The recommended approach in the cited account is the lowest effective energy, brief-pulse stimulation, a nondominant unilateral current, with treatments spaced **2 to 5 days apart**.

The clinical reasoning is not "TBI excludes ECT," nor is it "ECT is routine." First confirm that the depressive target is severe and genuinely refractory, reconsider whether cognitive symptoms, apathy or medication adverse effects are being mislabelled as persistent depression, and ensure that the procedural plan permits close cognitive observation. The wider interval and conservative technique make the cognitive-safety principle explicit. Procedural consent, anaesthetic assessment and ordinary medical safeguards remain part of local practice, but no unsupported protocol is supplied here.

Repetitive transcranial magnetic stimulation requires caution after TBI because it may increase seizure risk. This is a warning against casual extrapolation from ordinary depression practice. It does not establish that the procedure can never be used, and it does not justify inventing a stimulation protocol. Where neurostimulation is considered, specialist assessment should integrate the injury history, current cognitive status, seizure-related risk and the strength of the indication.

GOOD TO REMEMBER

Previous traumatic brain injury is not, by itself, a prohibition against electroconvulsive therapy. The examination-quality answer couples that statement with conservative technique, wider spacing, cognitive observation and caution with repetitive transcranial magnetic stimulation.

12.5 Acute agitation versus chronic aggression

Acute post-TBI agitation requires a cognitive-protection strategy. Avoid anticholinergic medicines and benzodiazepines because they can worsen cognition and are deliriogenic, and be cautious with typical antipsychotics because cognitive recovery may be adversely affected. Before

adding a sedative drug, reassess whether pain, fear, overstimulation, an unfamiliar environment or an iatrogenic adverse effect is sustaining the behaviour. Medication that suppresses expression without correcting the driver may complicate neurological and psychiatric reassessment.

The evidence behind antipsychotic caution must be stated at the correct altitude. Animal work suggests that haloperidol and risperidone may impair cognitive recovery; in one rat study haloperidol, but not olanzapine, impaired cognitive performance after TBI. This is a safety signal, not a human comparative efficacy trial and not proof that every atypical antipsychotic is cognitively harmless. Also examine for akathisia, because antipsychotic-induced motor restlessness may be misread as worsening aggression and treated with further antipsychotic exposure.

Chronic post-TBI aggression has a different pharmacological frame. Options include amantadine, beta-blockers such as propranolol or pindolol, and SSRIs such as sertraline; anticonvulsants may be considered despite limited data. Antidepressant treatment alone may be insufficient. Beta-blockers more reliably suppress aggressive behaviour, while valproate and carbamazepine have been reported to suppress violent outbursts. Agent choice should follow the dominant target and the patient's medical and cognitive tolerability, not a reflex hierarchy detached from context.

The controlled evidence tempers enthusiasm. Trials using **amantadine 200 mg/day** and **carbamazepine up to 400 mg/day** for post-TBI irritability, aggression and anger found no significant active-placebo difference on most outcome measures. These are doses from negative trials, not target doses. The result does not prove that the named agents never help an individual patient; it means the clinician must be modest about expected benefit and explicit about the stopping rule.

> PITFALL
>
> Escalating antipsychotics for apparent “aggression” can worsen the very picture being treated when the behaviour is actually akathisia, cognitive toxicity or delirium. Re-examine the phenotype before increasing exposure.

12.6 Disinhibition, PTSD and anxiety

Post-TBI disinhibition may be approached with SSRIs or anticonvulsants used along the same broad lines as in mania. SSRIs may be useful when reduced sexual drive would also address hypersexual behaviour. In a comportment vignette, sertraline reduced irritability, outbursts and hypersexuality. The goal is behavioural control with preserved participation, not global emotional blunting. Treatment of coexisting cognitive impairment may itself improve behavioural control, reinforcing the need to combine medication with rehabilitation rather than treating disinhibition as an isolated chemical problem.

For trauma-related and anxiety presentations after TBI, SSRIs are a reasonable first-line pharmacological strategy because of their safety and tolerability. Prolonged exposure, cognitive processing therapy and cognitive behavioural therapy are favoured treatment modes, so the

prescription should sit inside a psychological formulation. Anticholinergic medicines and benzodiazepines should be avoided because their adverse cognitive effects are especially problematic after brain injury.

The largest placebo-controlled trial of prazosin for trauma-related nightmares was negative. The lesson is not that nightmares are unimportant, but that a familiar indication outside TBI should not be imported as settled evidence without qualification. Describe the target, prioritise the supported psychotherapeutic mode, and review whether a medication produces meaningful sleep and daytime benefit without cognitive cost.

Disinhibition, anxiety and agitation can look superficially similar when the history is compressed into "restless and difficult." Their pharmacological implications differ. Disinhibition is an impulse-control and comportment target; anxiety requires a fear-based formulation; acute agitation demands an immediate search for drivers and cognitive toxicity. The prescription follows that distinction.

12.7 Psychosis and secondary mania

For psychosis due to TBI, atypical antipsychotics are likely to remain the main pharmacological treatment by extrapolation from idiopathic psychosis, because direct clinical trials in TBI are absent. The phrase "by extrapolation" matters. It declares the evidence gap and prevents ordinary antipsychotic practice from being presented as TBI-specific proof. Choose and monitor treatment with particular attention to neurological and cognitive tolerability. Receptor pharmacology and general comparative drug selection belong in the Schizophrenia: management, antipsychotics and adverse effects notes.

Secondary mania after TBI may reasonably be managed along the broad lines of idiopathic bipolar disorder. Lithium and valproate are options, while case reports describe atypical antipsychotics such as quetiapine or olanzapine. Persistent cognitive impairment increases susceptibility to neurological adverse effects from lithium and anticonvulsants. This shifts the balance toward cautious initiation, close observation and a low threshold to reconsider treatment when neurological or functional decline emerges, without supplying unsupported thresholds or serum targets.

The general pharmacokinetics, therapeutic range and non-TBI adverse-effect teaching for lithium belong in the Mood disorders: pharmacotherapy notes; anticonvulsant pharmacology belongs in the Epilepsy and psychiatry notes. In this section the examinable point is the acquired-brain-injury modification: usual syndrome-based treatments remain available, but evidence is often extrapolated and neurological tolerability may be poorer.

When psychosis, mania and aggression coexist, avoid building a drug for every symptom. Decide which syndrome best organises the presentation and whether a single agent can address the dominant target. Polypharmacy makes attribution difficult: improvement cannot be assigned confidently, and cognitive or neurological worsening may be hidden among overlapping adverse effects.

12.8 Apathy is not simply untreated depression

Post-TBI apathy calls for a different pharmacological hypothesis from depression. Case reports and case series support trials of prodopaminergic agents such as methylphenidate, amantadine or bromocriptine, while bupropion is generally avoided because it lowers the seizure threshold in a population already at raised seizure risk. The evidence level is modest, so treatment should be framed as an individual trial with a behavioural and functional endpoint, not as a guaranteed restorative treatment.

The key examination distinction is direction of effect. SSRIs are first-line pharmacotherapy for post-TBI depression, but studies in the general population suggest that SSRIs may worsen apathy. Therefore reduced initiation should not automatically trigger antidepressant escalation. Reassess whether the patient has sadness, distress and depressive cognitions, or instead a primary reduction in self-generated, goal-directed behaviour. The full phenomenology belongs with the psychiatric sequelae of traumatic brain injury; the prescribing consequence belongs here.

A useful endpoint is not that the patient appears more animated during a brief interview. Look for initiation of ordinary activity, participation in rehabilitation, completion of routines and reduced need for prompting. Family observations can help, but the target should be concrete enough that expectancy and general activation are not mistaken for meaningful recovery. Cognitive treatment may also improve behavioural control and engagement, so a medication trial must remain embedded in rehabilitation.

Some agents appear in more than one indication in this topic. Methylphenidate may have relevance to depression, cognition, fatigue and apathy; amantadine appears among options for apathy and chronic aggression. This overlap is not a licence for nonspecific prescribing. It is a reason to state the principal target before treatment and to judge the outcome against that target.

12.9 Review, adequacy and rational stopping

An adequate therapeutic trial balances caution with sufficiency. Begin at a tolerable intensity, progress carefully, and allow adequate exposure before concluding that a treatment has failed. At the same time, each addition should carry a review point and a reason to stop. This prevents “go slow” from becoming passive under-treatment and prevents partial response from becoming permanent polypharmacy.

At review, ask linked questions. Did the target improve in a way that matters? Was cognition, arousal, neurological function or rehabilitation participation worsened? Is the medicine still necessary in the present phase and locus of care? A drug that was defensible during inpatient containment may no longer be appropriate after the environment settles. Conversely, a long-term affective or behavioural syndrome may require a sustained plan rather than repeated short, poorly evaluated trials.

Apparent non-response should trigger reformulation before augmentation. Check that the original target was correctly identified, that the patient received an adequate trial, that adverse effects are not mimicking the disorder, and that psychosocial or environmental drivers are being

treated. When benefit is absent or outweighed by cognitive cost, taper according to current agent-specific guidance rather than leaving the medicine in place "because of TBI."

> After TBI, prescribe for one explicit target, protect cognition and seizure safety, test benefit against function, and stop medicines that cannot justify their place.

The mature examination answer is therefore organised, not encyclopaedic. Start with locus, focus and modus; state nonpharmacological priority; distinguish acute agitation from chronic aggression; give syndrome-specific choices with their evidence limits; and close with monitoring and deprescribing. Doses should be printed only when TBI-specific evidence supplies them and should retain their trial or vignette context. That discipline is clinically safer than a remembered formulary and academically more honest than false precision.

Delirium and the amnestic syndromes

13 · Delirium: definition, diagnostic criteria and clinical subtypes

The marks lie in recognising a syndrome, not in calling every confused patient delirious. **Delirium** is an acutely developing, fluctuating disturbance centred on attention and awareness, accompanied by disturbance in another cognitive domain and attributable to a physiological cause. DSM-5-TR classifies it as a neurocognitive disorder, but the bedside task is broader than attaching a label: demonstrate change from baseline, identify fluctuation, establish the cognitive pattern, respect exclusions and connect the syndrome to a medical, toxic, medication-related or withdrawal process.

Several distinctions prevent a vague answer. Attention is not the same as orientation, although each may be impaired. The syndrome is also not identical to the pathological process producing it: **acute encephalopathy** refers to the rapidly developing underlying brain process, whereas delirium names its clinical syndrome. The major diagnostic systems overlap in clinical description but organise the diagnosis differently. A disciplined answer therefore defines the syndrome, states each framework on its own terms, and then classifies psychomotor activity and course.

0.4%	6%-56%	UP TO 70%	UP TO 80%
COMMUNITY ADULTS	INCIDENCE IN HOSPITAL	MECHANICAL VENTILATION	TERMINAL ILLNESS NEAR DEATH

13.1 The clinical definition: a pattern of change

The defining disturbance is in **attention**: the capacity to direct, focus, sustain and shift mental engagement is reduced. This explains why a patient may answer an opening question yet lose the thread, become captured by irrelevant stimuli, or fail to follow a brief exchange. Attention must be distinguished from the amount of information remembered. Poor registration may follow poor attention, but an isolated memory complaint does not by itself establish delirium.

Awareness concerns orientation to and engagement with the environment. Its disturbance makes the patient less accurately situated in the immediate context. The diagnostic emphasis is on a departure from the person's usual state. A low cognitive baseline does not prevent delirium; it makes collateral comparison more important. The clinician asks what changed, when it changed, and whether the degree of engagement varies across observations.

The temporal signature is acute development with fluctuation. Under DSM-5-TR, the change usually emerges over **hours to a few days** and tends to vary in severity during the course of a day. A lucid interview therefore cannot erase convincing evidence of impairment earlier in the same illness. Conversely, a single disorganised interaction is insufficient unless the full pattern and physiological attribution are demonstrated.

Another cognitive disturbance completes the clinical picture. The affected domain may involve memory, orientation, language, visuospatial function or perception. Abnormal perceptions are common, especially visual or tactile hallucinations, but they are accompanying phenomena rather than the diagnostic centre. Behavioural or emotional disturbance may be prominent; neither substitutes for the required disturbance of attention and awareness.

RULE **Diagnose the pattern of acute change, not the loudest symptom.** Impaired attention and awareness, fluctuation, additional cognitive disturbance and a physiological basis carry more diagnostic weight than agitation, hallucinations or disorientation considered alone.

13.2 DSM diagnostic criteria: retain the full sequence

DSM-5-TR presents the diagnosis as an ordered criteria set. The order is educationally useful: define the core disturbance, establish its time course, add another cognitive domain, apply the neurocognitive and arousal exclusions, and require evidence of physiological causation. Omitting the final steps turns a syndrome definition into a symptom list.

CRITERION	REQUIRED CONTENT	DIAGNOSTIC WORK PERFORMED
A	Disturbance of attention accompanied by reduced awareness of the environment.	Identifies the core cognitive state rather than an isolated behavioural symptom.
B	Development over a short period, change from baseline, and fluctuation in severity during the course of a day.	Establishes onset, comparison state and variability.
C	An additional disturbance involving memory, disorientation, language, visuospatial ability or perception.	Shows that the syndrome extends beyond attention and awareness.
D	The disturbances are not better explained by a pre-existing, established or evolving neurocognitive disorder, and do not occur in severely reduced arousal such as coma.	Prevents misclassification of another neurocognitive state or coma.
E	History, examination or laboratory evidence supports direct physiological causation by a medical condition, intoxication, withdrawal, medication, toxin or multiple causes.	Links phenomenology to a physiological aetiology.

Criterion B requires the whole temporal sequence to remain visible: brief development, change from baseline, and fluctuation. “Acute confusion” is too imprecise because it does not show what changed or whether the course varies. Criterion C should also be stated as an additional cognitive disturbance, not as mandatory disorientation. A patient need not display every listed cognitive manifestation.

Physiological attribution is not an optional afterthought. The evidence may come from the history, physical examination or laboratory findings and may support one cause or several. The diagnostic answer should therefore pair syndrome recognition with a causal formulation, while the detailed aetiological model and pathophysiology remain separate topics.

MNEMONIC

A-B-C-D-E

Attention and awareness; **B**rief development, baseline change and fluctuation; **C**ognition additionally disturbed; **D**ifferential and severely reduced arousal exclusions; **E**tiological evidence. This preserves the diagnostic sequence without replacing the formal wording.

13.3 The highest-yield exclusion: another neurocognitive disorder or coma

The most commonly truncated part of the DSM-5-TR answer is Criterion D. The disturbance must not be better explained by another pre-existing, established or evolving neurocognitive disorder. This does not mean that a person with an established neurocognitive disorder cannot

also develop delirium. It means that the new attentional and awareness disturbance, its time course and its fluctuation must require an explanation beyond the prior condition.

Baseline is therefore a clinical comparator, not a screening score. Collateral history should establish the person's usual attention, awareness, communication, cognition and function, then describe the new departure. If previous impairment was already present, the clinician must demonstrate an acute superimposed change rather than simply relabel chronic cognitive difficulty. The full dementia syndrome and its workup belong in the Dementias and neurocognitive disorders notes.

The same criterion states that the disturbances do not occur in the context of a severely reduced level of arousal such as coma. This boundary matters because assessment of attention and awareness requires sufficient arousal for meaningful engagement. A patient who cannot be aroused enough to participate is not brought within the DSM-5-TR delirium criteria merely because a severe physiological brain disturbance is present.

GOOD TO REMEMBER

Write the baseline and arousal clauses explicitly. “Acute onset, fluctuating course and inattention” is recognisable shorthand, but it leaves the most discriminating exclusion unstated. In a vignette, ask whether the observed change exceeds the previous cognitive state and whether arousal permits attention to be assessed.

Disorientation, hallucinations, sleep-wake disturbance, emotional change and motor behaviour may strengthen the phenomenological description, but none repairs a missing Criterion D analysis. Similarly, evidence of infection, metabolic disturbance or drug exposure does not prove delirium unless the clinical syndrome is present. The diagnosis requires the appropriate phenomenology together with physiological attribution.

13.4 ICD characterization and etiological coding

ICD-11 characterises delirium through disturbed attention and awareness that develops over a short period and tends to fluctuate during the course of a day, accompanied by other cognitive impairment. The additional impairment may involve memory, orientation, language, visuospatial ability or perception. Behavioural and emotional disturbances are often present, and intoxication or withdrawal may cause the syndrome.

This characterization is clinically close to the shared syndrome core, but its classification spine is etiological. ICD-11 begins with **6D70** Delirium and then records whether the cause is a disease classified elsewhere, a psychoactive substance including medication, multiple etiological factors, another specified cause, or an unknown or unspecified cause.

ICD CATEGORY	ETIOLOGICAL MEANING	HOW TO USE IT CONCEPTUALLY
6D70.0	Delirium due to disease classified elsewhere	Name the syndrome and link it to the responsible disease.
6D70.1	Delirium due to psychoactive substances, including medications	Keep exposure, intoxication, withdrawal and medication context explicit.
6D70.2	Delirium due to multiple etiological factors	Avoid forcing a single-cause label when several factors contribute.
6D70.Y	Other specified delirium	Use when the cause is specified but not represented by the preceding categories.
6D70.Z	Delirium of unknown or unspecified cause	Preserve diagnostic uncertainty rather than inventing attribution.

Substance-class coding can become more specific, including alcohol-induced and sedative, hypnotic or anxiolytic-induced delirium. Those code branches establish aetiological placement; they do not replace a full withdrawal assessment. Withdrawal timelines, severity instruments and treatment regimens belong in the Substance use disorders notes.

13.5 The diagnostic systems must not be fused

The safe comparison begins by stating the overlap. Each framework centres disturbed attention and awareness, short development, fluctuation and additional cognitive impairment. Each permits physiological causes that include medical illness and substance-related states. This common clinical core supports recognition, but it does not authorise the candidate to build a hybrid list and attribute it to either manual.

TRAP **Do not present one harmonised "DSM and ICD criteria" list.** DSM-5-TR explicitly places the "not better explained" and severely reduced arousal or coma exclusion in Criterion D, whereas ICD-11 supplies the verified clinical characterization and organises the classification primarily by aetiology. State the frameworks separately, then compare them.

The distinction is more than examination etiquette. A formal exclusion determines whether a patient belongs inside a diagnostic set; an aetiological branch determines how an established syndrome is classified by cause. Treating these as interchangeable erases the work each system is doing. In a structured answer, lead with ICD-11, cross-map to DSM-5-TR, and identify the coma or severely reduced arousal clause as specific to the latter criteria presented here.

- For DSM-5-TR, preserve the complete A-to-E diagnostic logic and make Criterion D visible.
- For ICD-11, state the verified characterization without reconstructing an unverified enumerated essential-features list.
- For ICD-11 coding, place aetiology at the centre and retain an unspecified category when causation remains uncertain.

- In either system, do not let agitation, hallucinations or disorientation displace attention, awareness and temporal change.

This method also prevents false certainty. A patient can satisfy the clinical syndrome while the exact cause remains under investigation. Conversely, a plausible physiological insult does not establish delirium without the characteristic cognitive and temporal pattern. Syndrome recognition and causal attribution are linked tasks, but they are not the same inference.

13.6 Psychomotor subtypes: hyperactive, hypoactive and mixed

Psychomotor subtyping describes the activity pattern within established delirium. It must not become a substitute for the diagnostic criteria. The activity level can guide recognition and communication, but attention and awareness remain disturbed across subtypes. A quiet presentation is therefore fully compatible with delirium.

Hyperactive delirium has increased psychomotor activity. Agitation, mood lability and refusal to cooperate with medical care may accompany it; hallucinations and inappropriate behaviour may also be evident. It is more frequently recognised and is often associated with medication side effects and drug withdrawal. Its visibility explains why clinicians may mistake motor activation for the syndrome itself.

Hypoactive delirium has reduced psychomotor activity, with sluggishness and lethargy that may approach stupor. It is more frequent in older adults and is often unrecognised in emergency departments and hospitals. The Maudsley account associates this subtype with poorer prognosis. The central clinical danger is omission: reduced movement may be read as sleepiness, fatigue, low mood or cooperative calm while the attentional disturbance passes unnoticed.

Mixed level of activity includes patients whose activity rapidly fluctuates and also those with an apparently normal psychomotor level despite disturbed attention and awareness. The Maudsley description expresses the mixed pattern as features of increased and reduced motor activity. These descriptions are complementary reminders that activity can alternate, and that normal-looking motor behaviour does not exclude the cognitive syndrome.

> **PITFALL**
>
> Do not screen only the restless patient. Hyperactivity announces itself; hypoactivity can disappear into the ward routine. When behaviour becomes unusually quiet, slow or disengaged, assess attention, awareness, baseline change and fluctuation rather than assuming that absence of agitation means absence of delirium.

13.7 Course specifiers and the meaning of persistence

DSM-5-TR separates course from psychomotor activity. **Acute course** lasts **a few hours or days**, whereas **persistent course** lasts **weeks or months**. In hospital settings, delirium usually lasts about one week, although some symptoms often continue after discharge.

Course specifiers answer a different question from onset. The syndrome may begin over a short period yet continue beyond the immediate precipitating event. Persistence therefore does not by

itself convert delirium into a chronic neurocognitive disorder, nor does prolonged cognitive difficulty permit the clinician to abandon review of baseline, fluctuation, arousal and physiological attribution.

Activity subtype and duration should be documented independently. A patient may be predominantly quiet during a brief episode, activated during a prolonged episode, or alternate motor states while the syndrome continues. Collapsing these axes into labels such as “acute hyperactive delirium” without separately observing cognition, course and cause makes later reassessment difficult.

Documentation should keep the axes separate. Describe onset as the interval over which attention and awareness changed from baseline, and fluctuation as variation in severity across observation, collateral accounts and repeated examination. Record duration using the appropriate manual specifier. Record activity as hyperactive, hypoactive or mixed rather than inferring a stable subtype from a single encounter. This format lets another clinician revise one axis without discarding the entire formulation.

Persistent symptoms after discharge also make handover language important. The receiving clinician needs the previous baseline, the observed syndrome, the suspected cause, the pattern of fluctuation and the residual deficits. A bare statement that the patient “remains confused” loses the diagnostic architecture and encourages chronicity to be mistaken for a different syndrome without reassessment.

13.8 Bedside synthesis and examination answer

The clinical synthesis moves from observation to formulation. Establish that arousal permits meaningful assessment. Demonstrate impaired attention and reduced environmental awareness. Obtain collateral evidence of acute change and fluctuation. Identify another cognitive disturbance. Decide whether a pre-existing or evolving neurocognitive disorder better explains the presentation. Then seek evidence for a direct physiological cause and classify activity and course.

Bedside screening instruments can operationalise parts of this assessment, but they do not replace the diagnostic formulation; their administration and interpretation are covered in the dedicated delirium-screening section. Standard cognitive instruments are similarly adjuncts, not substitutes for collateral baseline and repeated observation, and their content and scoring belong in the Psychotherapies, clinical assessment, MSE and nosology notes.

The epidemiological gradient explains why case finding matters. Delirium is uncommon among community adults in the population estimate presented here, but it becomes much more frequent among medically vulnerable groups and in high-acuity settings. About **one-third** of patients older than 70 years admitted to hospital have delirium. Prevalence reaches up to 60% among nursing-home residents aged 75 years and older. These figures should sharpen vigilance without becoming a shortcut: setting changes prior probability, not the defining criteria.

A compact long answer begins by defining delirium through attention, awareness, acute baseline change, fluctuation, additional cognition and physiological attribution. It then states the ICD-11 characterization and etiological structure before cross-mapping the complete DSM-5-TR criteria and highlighting Criterion D. Psychomotor activity and duration are classified on separate axes. The answer closes by distinguishing delirium from acute encephalopathy, another neurocognitive disorder and severely reduced arousal. This sequence is concise without sacrificing the framework-specific distinction.

The answer should not drift into the separate territories of screening algorithms, aetiological mechanisms, investigation, management, prevention or withdrawal treatment. Name those tasks only when they clarify the boundary of the diagnosis. The present item is complete when the reader can recognise the syndrome, state the systems without conflation, apply the exclusion correctly and subtype the observed activity and course.

> Delirium is an acute, fluctuating disturbance of attention and awareness with additional cognitive impairment and physiological attribution; preserve the DSM-specific coma exclusion and the ICD etiological spine.

14 · The Confusion Assessment Method and delirium screening at the bedside

Bedside delirium screening converts a vague impression of "confusion" into an observable clinical argument. The marks lie in showing how change from baseline, fluctuation, attention, thought organisation and arousal are assembled, rather than merely expanding an acronym. A strong answer also separates detection from the search for cause: a screening result identifies a syndrome that needs clinical attention, but does not name the precipitating illness, replace physical assessment or prescribe management.

> **IN INDIA**
>
> Use the Confusion Assessment Method or another brief bedside screen, and train ward nurses and residents to apply it, because delirium is otherwise easily missed.

The **Confusion Assessment Method** is a brief, validated diagnostic algorithm. Its value is practical: it gives the examiner a repeatable structure while preserving the central role of history, collateral information, observation and mental status examination. The diagnostic criteria and clinical subtypes of delirium are considered separately in this chapter. Here the task is narrower and highly examinable: perform the screen correctly, apply its logical rule, document the evidence and recognise the limits of the result.

BASELINE	OBSERVE	TEST	INTEGRATE
ESTABLISH CHANGE	SEEK FLUCTUATION	ATTENTION AND AROUSAL	APPLY THE ALGORITHM

14.1 What screening is meant to accomplish

Screening should answer a focused question: does the present bedside pattern justify identifying delirium and pursuing the clinical consequences of that finding? It is not a generic cognitive

examination and it is not a substitute for formulation. Begin by defining the patient's usual cognitive and functional state, then establish what has changed. A patient who cannot supply a coherent history may still be assessed because the evidence comes from several sources: the patient's performance, observed behaviour, nursing observations, records and a reliable informant.

The screen must remain connected to clinical context. Reduced participation may arise from impaired arousal, language difficulty, sensory limitation, motor impairment, severe distress or unfamiliarity with the task. These possibilities do not make screening useless; they determine which observations are interpretable. Record adaptations and state when a feature cannot be assessed confidently. A forced binary answer conceals uncertainty and can be less useful than a transparent account of what was observed.

Before testing, make the encounter examinable. Introduce yourself, obtain the patient's attention, reduce avoidable distraction and use language the patient can understand. Check whether a response can be expressed by speech, gesture or another reliable mode. The screen should challenge cognition, not accidentally challenge hearing, vision, literacy or motor output. When a barrier persists, name it beside the affected feature. This prevents later readers from mistaking "not demonstrated" for "demonstrated absent."

- Establish the usual baseline and the nature of the recent change.
- Use collateral and ward observations to reconstruct variation across the clinical course.
- Test attention directly rather than inferring it from social fluency.
- Describe the evidence for each algorithmic feature before recording the overall result.

RULE **A screen is an organised bedside inference, not a score detached from its observations.** Preserve the chain from baseline and behaviour to feature-level findings, algorithm and clinical interpretation.

14.2 The clinical feature set of the CAM

The instrument primarily concentrates on **four** major clinical features. **Acute onset** asks whether the present state represents a recent departure from the patient's usual condition. **Fluctuation** asks whether the manifestations vary during the course rather than remaining uniform. These are temporal observations, so the bedside interview alone is often insufficient. The examiner must compare the current encounter with earlier observations and with the patient's baseline.

Inattention concerns failure to direct, sustain or manipulate attention adequately for the task. It should be tested and also observed throughout the interaction. **Disorganised thinking** is reflected in a failure to maintain a coherent, relevant or logically connected line of thought. **Altered level of consciousness** concerns arousal rather than the content of thought. Thought disorder and altered arousal are therefore alternative downstream features in the diagnostic rule, not synonyms for poor attention.

The full instrument also looks for disorientation, memory impairment, psychomotor change, disturbance of the sleep-wake cycle and perceptual disturbance. These observations enrich the

clinical description but should not be substituted for the defining algorithm. This distinction matters in written answers: describing many compatible symptoms is not equivalent to demonstrating that the required logical combination is present.

ASSESSMENT AXIS	BEDSIDE QUESTION	BEST SOURCE OF EVIDENCE
Onset	Is this a recent change from the usual state?	Collateral history, records and prior observation
Course	Has the presentation varied across the clinical course?	Serial examination and the ward narrative
Attention	Can the patient engage, sustain and manipulate a simple mental set?	Direct task performance plus behaviour across the interview
Thinking	Are answers relevant, coherent and logically connected?	Conversation and responses to structured questions
Consciousness	Is the level of arousal altered?	Direct observation before and during testing

14.3 Applying the diagnostic algorithm

The CAM diagnostic rule has a conjunctive core and an alternative final limb. The patient must have evidence of acute onset, a fluctuating course and inattention. In addition, the examiner must identify either disorganised thinking or an altered level of consciousness. The word “either” is important: both final-limb features may be present, but the algorithm does not require both. Conversely, a final-limb feature cannot compensate for absence of the temporal and attentional core.

This logical architecture prevents two common shortcuts. The first is calling any disorientation delirium without establishing change, fluctuation and attention. The second is excluding delirium because the patient is not visibly agitated or grossly incoherent. Neither shortcut applies the algorithm. The examiner should instead move feature by feature, decide whether the evidence is present, absent or not interpretable, and only then integrate the result.

MNEMONIC

AFA + D/C

Acute onset, **F**luctuating course and **A**ttention failure form the required core; add either **D**isorganised thinking or altered **C**onsciousness. The letters preserve the AND-plus-OR structure rather than creating a simple symptom list.

The result should be documented as a reasoned conclusion. A useful note states the baseline source, the observed change, evidence for variability, the attention task and behaviour, the organisation of thought, the level of arousal, and the resulting algorithmic interpretation. If collateral information is unavailable, say so. If language or motor limitation prevents a task, describe the limitation and use other interpretable evidence without pretending that an unperformed test was normal.

TRAP **Do not reduce the CAM to "any feature is positive."** The temporal change and fluctuation, inattention, and the final either-or limb have distinct logical roles; a list of compatible symptoms does not reproduce the algorithm.

14.4 Establishing onset and fluctuation at the bedside

Acute onset is relational: it describes change from the patient's own baseline, not simply poor performance compared with an ideal norm. Ask an informant what the patient could usually do, communicate and understand before the present illness. Establish when the difference was first recognised and what made it noticeable. Records and nursing notes may show that the patient's engagement, speech or arousal differs from an earlier encounter even when the exact beginning cannot be reconstructed.

Fluctuation is similarly longitudinal. A calm, coherent interview does not erase a convincing history of marked variation, and a difficult encounter should not be treated as the whole course. Seek concrete contrasts: periods of better and poorer engagement, varying ability to follow conversation, or changes in arousal noted by staff or family. Avoid vague collateral such as "sometimes confused" when the informant can describe what the patient actually did and how that differed from usual functioning.

Operationally, ask for the last clearly usual state and the first clearly changed state without inventing precision that the history cannot support. Compare observations made in different care contexts and under different levels of fatigue or stimulation. Separate fluctuation in cognition or arousal from ordinary variation in cooperation, frustration or preference. When the bedside picture and the longitudinal account disagree, repeat the screen and preserve both observations for comparison.

GOOD TO REMEMBER

Before asking the patient to perform a cognitive task, ask the ward team what changed and how the state has varied. The temporal features may already be visible in the clinical narrative, while a single polished interview can conceal them.

14.5 Testing attention and observing arousal

Attention testing should be brief, comprehensible and matched to the patient's ability to hear, speak and respond. Common bedside operations include digit span, with repetition of **five digits forward** and **three digits backward**, and asking the patient to recite the days of the week or months of the year backward. These are not interchangeable with casual conversation. Familiar social responses can remain superficially fluent even when the patient cannot hold a simple mental set.

Observe the process, not merely completion. Does the patient register the instruction, lose the sequence, become captured by environmental stimuli, need repeated redirection or abandon the task? A failure may support inattention, but interpretation still requires checking comprehension and response capacity. Simplify the instruction when language is uncertain, ensure that hearing

and vision are adequate for the selected task, and document why a result is unreliable when these conditions cannot be met.

Arousal should be assessed before cognitive demand is increased. Note whether the patient can be engaged and maintain wakeful contact long enough for meaningful testing. Altered consciousness may itself satisfy the alternative final limb of the algorithm, but it also affects how much confidence can be placed in other responses. The examiner should not manufacture a detailed cognitive profile from a patient whose arousal prevents valid participation.

The Mini-Mental State Examination, Mini-Cog and Montreal Cognitive Assessment are commonly used bedside cognitive tests in delirium screening, but a tool does not replace the mental status examination. Their item content, scoring and standardised interpretation belong to the Psychotherapies, clinical assessment, MSE and nosology notes. In this context they may contribute cognitive observations; they do not replace the CAM logic or the temporal history.

14.6 CAM-ICU: validity with the population kept attached

The **CAM-ICU** is the critical-care application of the approach. In mechanically ventilated older patients, it showed **95% sensitivity** and **88% specificity** when compared with expert clinical diagnosis. Administration usually takes about 2 minutes. These figures make the instrument easy to remember, but their population label is part of the fact, not a footnote.

Sensitivity describes how effectively the instrument identifies patients who have the target condition against the chosen reference assessment; specificity describes how effectively it excludes those who do not. Neither measure turns a screen into an etiological diagnosis. The result must still be read with the bedside circumstances, the adequacy of testing and the clinical course. A technically completed screen cannot repair an incorrect baseline history or an uninterpretable response.

> **PITFALL**
> Do not quote the critical-care validity figures as though they were universal performance characteristics for every ward patient. They belong to the mechanically ventilated older population in which that application was assessed.

The same caution applies when communicating results. State the instrument and setting, the features found, the response limitations and the action required from the clinical team. “Screen positive” is less useful than a feature-level account showing which observations drove the result. If the state later changes, repeat assessment and preserve the comparison rather than overwriting the earlier description.

Serial use is clinically useful only when the conditions of assessment remain visible. A different result may reflect genuine change, a better interval, altered arousal, improved communication or removal of a testing barrier. Therefore compare feature-level findings, not merely the final label. Handover should identify which observation changed and which remained stable. This makes repeated screening a record of the clinical course instead of a pile of disconnected results.

14.7 Choosing among bedside instruments

Instrument choice depends on the clinical setting, the purpose of assessment, the patient's capacity to participate and the training available to the examiner. The aim is not to collect overlapping scores. Choose a tool that answers the immediate detection question, then use the clinical examination to understand the result and its limitations.

INSTRUMENT	ESTABLISHED USE IN THE AVAILABLE EVIDENCE	BOUNDARY FOR BEDSIDE USE
Confusion Assessment Method	Validated in inpatient units, emergency rooms and intensive-care settings	Apply its feature logic; do not treat it as a cause-finding workup.
4A's Test (4AT)	A quick delirium screening tool that does not require special training to administer	Use the authorised instrument; its scored items and thresholds are not reproduced here.
CAM-ICU	Validated for critically ill patients in the intensive-care setting	Keep setting-specific evidence attached to any interpretation.
Intensive Care Delirium Screening Checklist	Validated for critically ill patients in the intensive-care setting	Do not import an unverified scoring rule from memory.
Delirium Rating Scale-Revised-98	Validated in general medical and geriatric populations	Use according to its authorised administration and interpretation.
Mini-Mental State Examination, Mini-Cog and Montreal Cognitive Assessment	Common bedside cognitive tests used during delirium screening	Cognitive screening supports assessment but does not replace the mental status examination.

The 4AT is attractive where a rapid tool without special training is needed, but its worksheet must not be reconstructed from recollection. The Intensive Care Delirium Screening Checklist and CAM-ICU belong to critical-care practice, while the Delirium Rating Scale-Revised-98 has validation in general medical and geriatric populations. Naming these settings is more informative than producing an ungrounded league table of scores.

14.8 Interpretation, documentation and examination synthesis

The final product is a clinical statement that another doctor can verify. Record the informant and baseline, the evidence for recent change and variation, the attention task used, the observed organisation of thought, the level of arousal, any barriers to assessment and the algorithmic conclusion. This format makes disagreement productive: a later examiner can identify whether the difference lies in the history, the observed feature or the rule's application.

A positive result should trigger the clinical response appropriate to suspected delirium, including investigation of the underlying cause and attention to immediate safety, but those management details belong in the dedicated investigation and management section. A negative result should not silence a strong longitudinal concern when the assessment occurred during a better interval

or when the patient could not be tested adequately. Re-examination is part of honest interpretation because the target phenomena can vary.

Common documentation failures are avoidable. Do not write “confused” without describing the change. Do not equate disorientation with the complete algorithm. Do not report an attention task as passed when the patient never understood the instruction. Do not quote a validity estimate without its population. Do not let a cognitive screening score replace observation of arousal, attention and course. Each error breaks a different link in the chain from bedside evidence to inference.

The wording of uncertainty matters. “Feature absent” means the available observations argue against it. “Feature not established” means the history or examination did not supply enough evidence. “Feature not assessable” means a barrier prevented a valid judgement. These statements lead to different next steps: accept the observation, obtain better collateral, adapt the examination or repeat it later. Collapsing them into a negative result creates false reassurance and weakens continuity between teams.

The concise examination answer therefore has a recognisable spine. Define the Confusion Assessment Method as a brief validated diagnostic algorithm. State and explain the core feature set, distinguishing temporal change, attention, thinking and consciousness. Apply the required AND-plus-OR rule rather than listing symptoms. Describe bedside attention testing and the need for collateral and serial observation. Name setting-appropriate alternatives without reproducing unsupported scoring systems. Finish with validity kept population-specific and with the screen’s clinical limits.

> CAM requires acute onset, fluctuation and inattention, plus either disorganised thinking or altered consciousness; document the evidence for every limb before declaring the result.

15 · Delirium: aetiology, predisposing and precipitating factors, and pathophysiology

The examiner is asking for a causal formulation, not a catalogue of diseases. A strong answer explains why the same medical insult produces delirium in one patient but not another, then connects that vulnerability model to disturbances of arousal and attention and cerebral regulation. The central argument is that delirium arises from an interaction between the patient’s baseline reserve and the burden of new insults. Aetiological categories identify what can supply those insults; pathophysiological hypotheses explain how very different causes can produce a common syndrome.

This section therefore moves from bedside causation to brain mechanism. It does not restate the diagnostic criteria, clinical subtypes, screening instruments, investigation, management, prevention, or outcomes, which are addressed in their respective sections. Substance-withdrawal states are included only as causes and as mechanistic contrasts; their timelines, severity assessment and treatment belong in the Substance use disorders notes.

BASELINE	ACUTE	NETWORK	EEG
VULNERABILITY	PRECIPITANT	DYSREGULATION	PHYSIOLOGY

15.1 The multifactorial vulnerability model

The **multifactorial vulnerability model** separates background susceptibility from the event that brings the syndrome into expression. Predisposing factors are characteristics already present before the acute episode; precipitating factors initiate or lead to delirium. The distinction is conceptual rather than merely chronological. Baseline vulnerability lowers the amount of additional physiological or environmental stress required, while a robust patient generally needs a larger or more cumulative burden of stressors.

This is often called a **two-hit** model, but the label should not be misread as a literal event count. It describes the interaction of vulnerability and insult. Either side may itself be cumulative: frailty, sensory impairment and chronic disease may coexist before admission, while infection, medication changes, sleep loss and immobilisation may accumulate afterwards. The explanatory value lies in the balance between the sides, not in counting isolated causes.

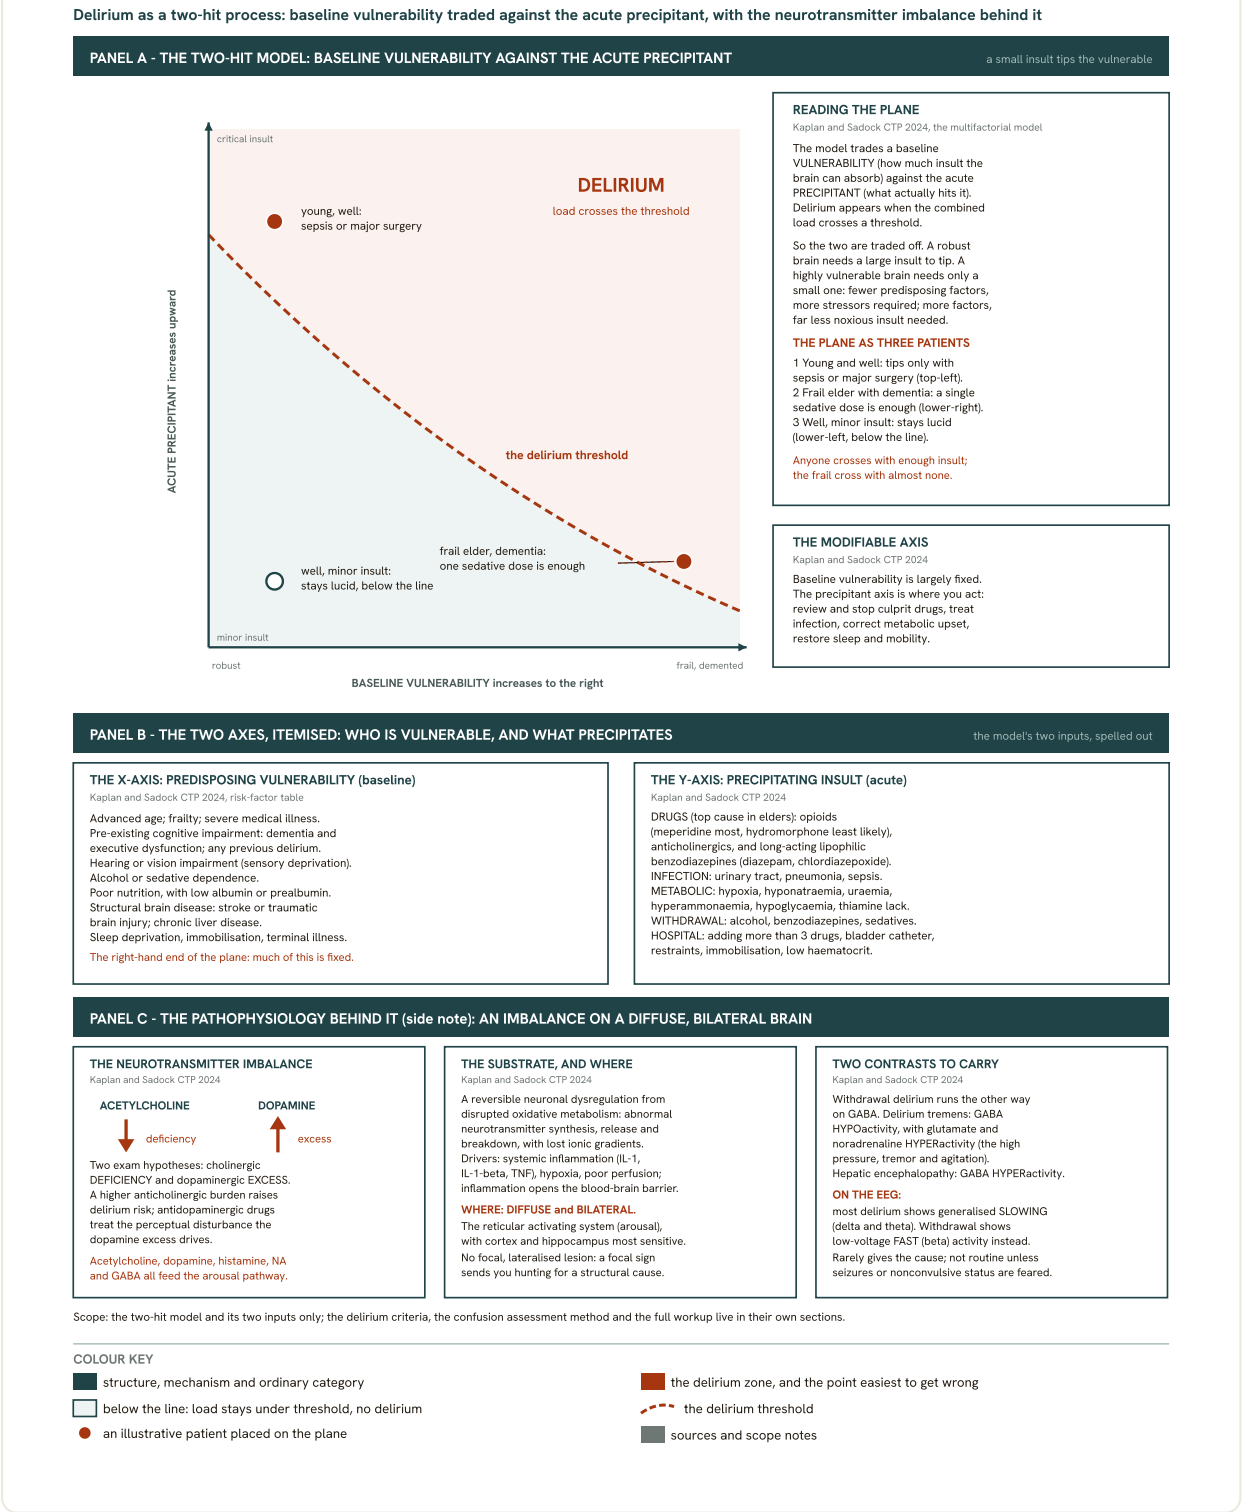


Fig 34.4 **Delirium as a two-hit process.** The multifactorial model plots baseline predisposing vulnerability against the acute precipitant: a robust brain crosses into delirium only with a large insult, while a highly vulnerable brain (advanced age, dementia, sensory impairment, frailty) needs only a small one, so the three example patients sit on either side of one threshold. Panel B itemises the two axes, and the side note gives the pathophysiology carried by these texts, a cholinergic deficiency and dopaminergic excess on a diffuse, bilateral brain, with the delirium tremens and hepatic encephalopathy GABA contrasts. (Kaplan and Sadock CTP 2024.)

The accompanying figure should be read across that balance. A highly vulnerable patient may cross the clinical threshold after a modest additional insult; a patient with little baseline vulnerability requires a greater acute burden. This explains why an apparently minor ward event

can be causally important in one person without being sufficient in another. It also explains why identifying one plausible precipitant should not end the causal search.

The model is probabilistic, not deterministic. A predisposing factor tells the clinician that the threshold may be lower; it does not prove that delirium must occur. A precipitant supplies causal pressure, but its significance depends on the host in whom it acts. This prevents symmetrical errors: dismissing a modest insult because it appears too small, or assuming that a large-looking insult explains the entire episode without considering the patient's prior vulnerability.

RULE **The smaller the baseline cerebral reserve, the smaller the insult needed to produce delirium.** Conversely, low baseline vulnerability generally requires a larger or cumulative precipitating burden.

15.2 Predisposing factors: the vulnerable host

Predisposing factors describe reduced reserve before the acute insult. The strongest clinical idea is not that each factor independently "causes" delirium, but that several factors can make cerebral function easier to destabilise. Pre-existing cognitive impairment, including dementia or executive dysfunction, and a previous history of delirium indicate an already vulnerable cognitive system. Advanced age, frailty and severe medical illness similarly signal limited capacity to absorb a new physiological disturbance.

Other vulnerabilities operate by reducing the reliability of input, physiological reserve or adaptation to the hospital environment. Hearing or visual impairment makes accurate interpretation of surroundings more difficult. Poor nutritional status, reflected by low albumin or prealbumin, chronic liver disease, terminal illness, sleep deprivation and immobilisation marks a host in whom metabolic and environmental stress may be less well tolerated. Alcohol or sedative dependence adds vulnerability because abrupt changes in exposure can become clinically important.

Structural brain disease, including stroke or traumatic brain injury, belongs on the predisposing side when it is established background pathology. The same diagnostic family may appear on the precipitating side when an acute brain event initiates the episode. Thus, classification depends on the temporal and causal role of the condition in the particular patient, not merely on its name.

Predisposition is also clinically dynamic. Sleep loss and immobilisation may already be present before an acute deterioration, then deepen during admission and function as ongoing precipitants. The useful question is therefore not "Which list owns this factor?" but "What role is it playing now?" Keeping baseline state and recent change in separate columns preserves the logic of the model even when the same exposure spans the transition into delirium.

- Cognitive reserve: pre-existing cognitive impairment and previous delirium.
- General reserve: advanced age, frailty, severe medical illness and terminal illness.
- Brain substrate: established stroke, traumatic brain injury or other structural brain disease.
- Sensory and nutritional reserve: hearing or visual impairment and poor nutritional status.

- Dependence and environment: alcohol or sedative dependence, sleep deprivation and immobilisation.

15.3 Precipitating factors: the acute burden

Precipitating factors are new insults that initiate delirium in a susceptible brain. Some are direct consequences of illness, while others are created or amplified by care. Bed rest or immobilisation, physical restraints, sleep deprivation and low haematocrit are recognised precipitants. Their importance is easily underestimated because none looks like a dramatic neurological event in isolation. Within the multifactorial model, however, several modest stresses may be sufficient when baseline vulnerability is high.

Medication exposure is particularly important. Opioids, anticholinergic medicines and benzodiazepines are prominent medication precipitants. In hospital, risk also rises with **more than three newly added medicines**. The count is not a diagnostic rule; it is a warning that rapid polypharmacy increases the chance of central effects, interaction, accumulation and attribution error. A medication list should therefore be read as a time sequence: what was introduced, stopped or escalated around the change in mental state?

Hospital processes may become part of the causal chain. Bladder catheters, malnutrition and iatrogenic events such as hospital-acquired infection or electrolyte abnormality are additional precipitants. A catheter, immobility, restraints, sleep loss and new medicines may accumulate during the same admission. Delirium may therefore emerge from a network of interacting insults rather than a single neat lesion.

A precipitating-factor formulation should preserve sequence and burden. It asks which exposure preceded the change, which persisted as the state evolved, and which appeared only after delirium was already established. This temporal discipline distinguishes plausible initiating factors from later consequences. It also keeps causal confidence proportional to the evidence: several contributors may be ranked as likely or possible without pretending that bedside chronology can always isolate a single sufficient cause.

QUESTION	PREDISPOSING SIDE	PRECIPITATING SIDE
When is it present?	Before the acute episode	At or near the onset
What does it do?	Sets baseline vulnerability	Initiates or leads to delirium
Typical examples	Cognitive impairment, frailty, sensory impairment, structural brain disease	New medicines, restraints, sleep deprivation, catheter use, iatrogenic events
Clinical reasoning	High vulnerability lowers the required insult	Low vulnerability requires a larger burden

PITFALL

Do not promote the first abnormality found into "the cause" without testing its timing and explanatory power. The multifactorial model permits several simultaneous precipitants, especially in a highly vulnerable inpatient.

15.4 Aetiological map: think beyond the ward chart

The **aetiological map** has broad families: systemic illness, medication effects, withdrawal states, primary brain disease, and surgery or trauma. These categories are not rival diagnoses. A single patient may occupy several columns, and the same process can affect the brain through more than one mechanism. The map prevents premature closure by forcing the answer to include systemic, toxic, withdrawal-related, neurological and procedural possibilities.

AETIOLOGICAL FAMILY	EXAMPLES WITHIN THE PERMITTED EVIDENCE	PATHOPHYSIOLOGICAL ROUTE TO CONSIDER
Systemic illness	Infection; dehydration, hypoxia, hypoglycaemia, hyperammonaemia, uraemia, hyponatraemia, thiamine deficiency; cardiovascular conditions	Inflammation, impaired oxygenation, perfusion or metabolism
Medicines	Anticholinergics, benzodiazepines and other sedative-hypnotics, opioids, antiarrhythmics, fluoroquinolones, clarithromycin, digoxin and interferons	Altered neurotransmission within arousal and attention systems
Withdrawal states	Alcohol, benzodiazepine or other sedative-hypnotic withdrawal	Loss of inhibitory tone with excitatory and noradrenergic overactivity in alcohol withdrawal
Primary brain disease	Stroke or transient ischaemic attack, trauma, meningitis, encephalitis, vasculitis, nonconvulsive status epilepticus and postictal states	Disruption of diffuse bilateral cortical and arousal networks
Surgery or trauma	Hip-fracture repair and coronary artery bypass grafting	Procedural stress superimposed on baseline vulnerability

MNEMONIC

System, Medicines, Withdrawal, Brain, Surgery

Use this sequence as a broad causal sweep: **S**ystemic illness, **M**edicines, **W**ithdrawal, primary **B**rain disease, and **S**urgery or trauma. It is a search scaffold, not a claim that the categories are mutually exclusive.

The category list also clarifies an important boundary. Nonconvulsive status epilepticus and postictal states can enter the delirium differential as primary brain causes, but seizure classification and anticonvulsant pharmacology belong in the Epilepsy and psychiatry notes.

Likewise, the presence of a withdrawal category does not authorise importing the full withdrawal syndrome into this causal section.

The categories converge rather than compete. Systemic illness may alter inflammation and perfusion; a medicine may alter neurotransmission; withdrawal may shift inhibitory and excitatory balance; and primary brain disease may directly disrupt distributed networks. Once several pathways converge on the same vulnerable arousal-attention system, the bedside phenotype may look similar despite very different starting points. That is why aetiology must be framed broadly while the mechanistic explanation remains integrated.

15.5 Older inpatients: infection and medication dominate

Among hospitalised older adults, infection and medication are the most important causes, and **medication is the single most common cause in older patients**. This emphasis follows directly from the vulnerability model. Older inpatients commonly bring baseline vulnerabilities into an environment where acute illness, procedures and medication changes accumulate. The aetiological answer should therefore give medication exposure the same seriousness as infection rather than treating it as an afterthought.

The important infectious causes in this setting include urinary tract infection, pneumonia and sepsis. The frequent medication causes are opioids, anticholinergic medicines and benzodiazepines. These groups also appear among general precipitants, but the older-inpatient emphasis is sharper: they are common enough that chronology and actual exposure deserve explicit reconstruction whenever mental state changes after admission.

Within opioids, the comparative signal is specific: meperidine is most likely and hydromorphone least likely to induce delirium. Within benzodiazepines, the most lipophilic agents with the longest elimination half-lives, including chlordiazepoxide, flurazepam and diazepam, are most likely to cause delirium. These comparisons should not be converted into a general prescribing algorithm here. Their purpose is aetiological: they refine suspicion when a delirium begins after exposure.

> **GOOD TO REMEMBER**
>
> In an older inpatient, reconstruct medication exposure as carefully as infection. Review new additions, cumulative sedative and anticholinergic burden, and whether several small hospital precipitants appeared together.

15.6 Neurotransmitter hypotheses: arousal and attention lose balance

The definitive pathophysiology remains incompletely defined, but neurotransmitter hypotheses provide a useful bridge from cause to phenotype. The central pairing is **cholinergic deficiency** with **dopaminergic excess**. Acetylcholine supports the distributed processing required for arousal and attention; reduced cholinergic function makes these systems less stable. Consistent with this model, greater anticholinergic burden increases delirium risk.

Increased dopaminergic neurotransmission is believed to contribute to the perceptual disturbances seen in delirium. The clinical observation that antidopaminergic antipsychotic

drugs can reduce such disturbances supports the hypothesis, but it does not establish dopamine excess as the sole cause of the syndrome. Drug selection, dosing and monitoring belong in the management section.

The arousal-attention pathway is chemically plural. Acetylcholine, dopamine, histamine, norepinephrine and GABA all influence it. Delirium should therefore not be described as a simple deficiency disease of one transmitter. Different aetiologies may perturb different parts of the system, while impaired metabolism, inflammation and altered perfusion can disturb several transmitters at once. The paired cholinergic and dopaminergic hypotheses are high-yield organising principles within a wider network model.

It helps to distinguish level of explanation. Neurotransmitter change is a downstream functional description, while infection, hypoxia, impaired perfusion or medication exposure may be upstream drivers. The same upstream insult may disturb several transmitter systems, and different upstream insults may converge on a similar transmitter imbalance. This layered account is more coherent than assigning every aetiology its own isolated pathway.

TRAP **Do not write "delirium is caused by acetylcholine deficiency."** Cholinergic deficiency and dopaminergic excess are implicated hypotheses within a broader disturbance involving multiple transmitters, metabolism, inflammation and distributed arousal networks.

15.7 Inflammation, oxidative metabolism and diffuse brain dysfunction

A second major account begins outside the neuron. **Systemic inflammation** is associated with increased inflammatory cytokines, including **IL-1 and IL-1-beta** and TNF. Inflammation can increase blood-brain-barrier permeability and affect microglia and astrocytes. The implication is not that every case is an inflammatory encephalitis, but that systemic illness can alter the cerebral milieu and destabilise neuronal regulation without requiring a focal lesion.

Other routes converge on energy failure. Hypoxia, impaired blood flow, reduced tissue perfusion and impaired metabolism can disrupt oxidative metabolism. Delirium is accordingly conceptualised as a reversible syndrome of neuronal dysregulation: energy-dependent ionic gradients become disturbed, and the synthesis, release or metabolism of neurotransmitters becomes abnormal. This framework links systemic infection, cardiopulmonary compromise and metabolic derangement to a common failure of neuronal organisation.

The neuroanatomical expression is **diffuse bilateral cortical dysfunction**, not a focal or lateralised process. The **reticular activating system** is assumed to be affected because it regulates alertness, while neocortical and limbic inputs to that system contribute to attention. The cerebral cortex and hippocampus are relatively more sensitive to metabolic disruption than other brain regions. Arousal, attention and memory can therefore fail together even when no single cortical territory accounts for the presentation.

This anatomical model has a practical interpretive consequence. A global change in mental state is compatible with diffuse metabolic or inflammatory dysfunction, whereas a focal neurological finding should prompt a search for a structural cause rather than being absorbed uncritically into

the label of delirium. Primary brain disease can cause delirium, but the clinical syndrome itself is not represented as a unilateral lesion.

The pathway can be summarised as a cascade without claiming a single final mechanism: systemic disturbance alters inflammation, oxygen delivery, perfusion or metabolism; the cerebral environment and cellular regulation become unstable; ionic gradients and transmitter handling are disrupted; and diffuse cortical plus arousal networks lose coordinated function. Reversibility is possible because the model concerns dysregulation rather than necessarily permanent neuronal destruction, although the underlying illness may still be serious.

15.8 State-specific mechanisms and the EEG correlate

Some aetiologies create recognisable transmitter contrasts. Delirium tremens is associated with GABA hypoactivity and hyperactivity of glutamate and norepinephrine. The increased noradrenergic drive of alcohol withdrawal helps account for high blood pressure, tremor and agitation. This is a mechanistic statement only; the withdrawal timeline, severity instruments and benzodiazepine regimens are taught in the Substance use disorders notes.

Hepatic encephalopathy provides the opposite GABA-direction contrast: GABA hyperactivity is implicated. Benzodiazepines show why transmitter labels must be applied in context. By enhancing GABA tone in corticothalamic networks, benzodiazepines can themselves produce cognitive impairment implicated in delirium, even though loss of inhibitory GABA activity characterises the withdrawal state. “More GABA” or “less GABA” is therefore not a universal delirium formula; the direction depends on the aetiological state and network involved.

The **electroencephalographic signature** of most delirium is generalised slow-wave activity in the delta and theta range, matching the concept of diffuse cerebral dysfunction. In hyperadrenergic alcohol or sedative withdrawal, the contrast is increased low-voltage fast beta activity. Thus, slowing is the usual exam answer, while fast activity points toward a withdrawal-related physiological state.

Electroencephalography rarely identifies the underlying cause and is not routinely performed solely because delirium is present. It becomes relevant when seizure activity or nonconvulsive status epilepticus is suspected. This preserves the distinction between a physiological correlate and an aetiological test: the tracing may support diffuse dysfunction or reveal a seizure-related alternative, but it does not replace the clinical causal formulation.

> Delirium is best understood as a vulnerable brain exposed to sufficient acute stress, producing reversible dysregulation of metabolism and neurotransmission across diffuse arousal and attention networks.

16 · Investigation and management of delirium

Delirium is a medical emergency with a psychiatric phenotype. The high-scoring answer therefore begins with a search for the precipitating illness or physiological derangement, not with a sedative prescription. Diagnosis remains clinical, but investigation explains why the brain has

become acutely dysfunctional and directs definitive treatment. The prescribing question comes later: whether dangerous agitation, severe distress or disruptive psychotic symptoms persist despite correction of causes and supportive care.

IN INDIA

Treat the underlying cause and use non-drug measures first; reserve haloperidol or olanzapine for dangerous agitation at the lowest effective dose, and keep dexmedetomidine for the intensive-care setting.

The practical hierarchy is **cause-directed treatment**, environmental and nursing support, then narrowly targeted symptomatic medication when the risk-benefit balance justifies it. Antipsychotics and benzodiazepines do not remove the cause. Formal diagnostic criteria, subtypes and bedside screening belong in their dedicated delirium sections; prevention bundles and longer-term outcomes are also addressed separately.

FIND	FIX	SUPPORT	RESERVE
THE PRECIPITANT	THE DERANGEMENT	THE PATIENT	SYMPTOM DRUGS

16.1 Management architecture: LOCUS, FOCUS and MODUS

LOCUS is determined by medical urgency and the resources needed for assessment, observation and treatment. A patient with delirium generally requires an emergency or inpatient medical setting, because the syndrome may signal an unstable underlying disorder and because cognition, cooperation and risk can change rapidly. Psychiatry contributes phenomenological assessment, behavioural containment and medication advice, but the patient remains under active medical evaluation rather than being diverted into a purely psychiatric pathway.

FOCUS changes across the episode. Acutely, priorities are physiological stability, identification of the precipitant, immediate safety and relief of severe distress. Once the cause is being corrected, focus shifts to restoring orientation, sleep-wake organisation, hydration, mobility and ordinary communication. As the state settles, the team reviews whether supportive measures can be simplified and whether any symptom-control medicine can be withdrawn. The diagnostic and treatment plan should be revisited whenever the course does not fit the presumed cause.

MODUS is deliberately multimodal. It combines investigation, correction of medical and medication-related contributors, nursing and environmental measures, family participation, mobilisation, sensory support and judicious pharmacotherapy. Psychological work in the acute state is supportive rather than insight-oriented: brief reassurance, repeated explanation and reduction of frightening misinterpretation. Social work may be needed to establish baseline function, identify available caregivers and make the discharge environment workable.

RULE **Treat the cause before treating the behaviour.** Symptom control may make care possible, but it is an adjunct to investigation and definitive medical treatment, never a substitute for them.

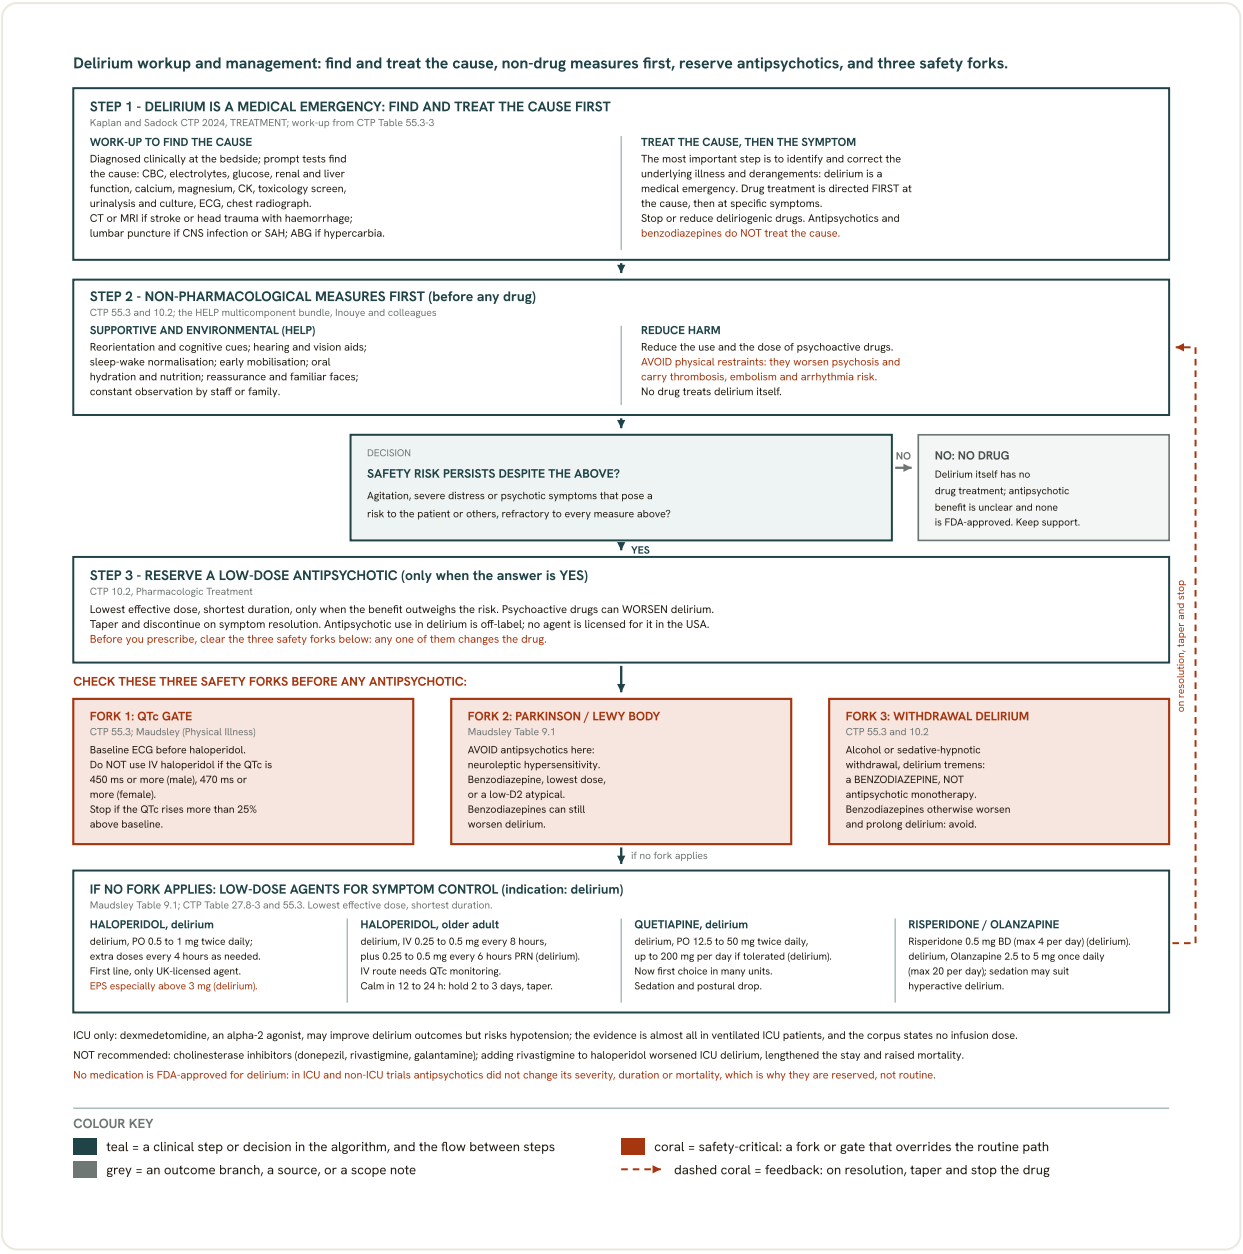


Fig 34.5 Delirium: the workup-and-management algorithm. Because delirium is a medical emergency, the path starts with finding and treating the cause and with non-pharmacological measures, not with a prescription; an antipsychotic is reserved for agitation, severe distress or psychosis that poses a safety risk and is refractory to those measures, at the lowest effective dose with a baseline ECG and QTc limits, and tapered on resolution. Three safety forks override the routine path: a baseline QTc gate before intravenous haloperidol; in Parkinson's disease and Lewy body dementia antipsychotics are avoided, and if a drug is essential a cautious lowest-dose benzodiazepine or a low-dopamine-blockade atypical such as quetiapine or clozapine is used under specialist guidance, recognising that benzodiazepines can themselves worsen delirium; and withdrawal delirium including delirium tremens is treated with a benzodiazepine rather than antipsychotic monotherapy. (Kaplan and Sadock 2024; Maudsley Prescribing Guidelines.)

The accompanying figure should be read as a recurrent loop rather than a single pass: stabilise, investigate, intervene, observe response and reopen the differential if the patient worsens or fails to improve. This protects against the seductive error of stopping the workup once a plausible contributor has been found.

16.2 Bedside investigation begins before the laboratory form

Investigation starts with verification of the syndrome and reconstruction of change from baseline. Establish what the patient was like before the acute illness, when the change was first

noticed, how alertness and attention have varied, and whether the presentation is compatible with delirium rather than a stable psychiatric syndrome. Obtain collateral history because the patient's account may be incomplete or internally inconsistent. The predisposing-versus-precipitating model is discussed in the dedicated aetiology section; here its purpose is to keep the search broad and prioritised.

The medication history is part of the examination, not an administrative appendix. Reconcile prescribed drugs, recent additions or withdrawals, non-prescription preparations, substances and possible access to another person's medication. Ask the ward team or family about actual exposure rather than relying only on the chart. Link each possible contributor to timing, physiological plausibility and change after correction. Avoid the opposite errors of ignoring medication effects and blaming every medicine without evidence.

Physical examination and observation guide test selection. Look for instability, dehydration, pain, infection, focal neurological findings, respiratory compromise, trauma and evidence of intoxication or withdrawal. Observe swallowing, mobility, continence, sensory impairment and the patient's ability to cooperate safely with essential care. A fluctuating behavioural presentation still requires an ordinary medical examination. Agitation is not evidence that examination can be deferred, while quietness is not evidence that the state is benign.

- Define baseline cognition, function and usual behaviour from a reliable informant.
- Reconstruct the temporal relationship between illness, procedures, medication changes and altered mental state.
- Repeat focused examination when the state changes rather than treating the first assessment as permanent.
- Document what remains unexplained, so diagnostic uncertainty stays visible to the next team.

PITFALL

Finding a urinary abnormality, sedating medicine or recent procedure does not prove that it is the sole precipitant. Delirium is often clinically multifactorial. Continue the search when severity, timing or response is not adequately explained.

16.3 Core laboratory and physiological workup

The **core delirium workup** is broad because reversible systemic disturbances can converge on a similar bedside syndrome. Tests should be interpreted in the clinical context, not ordered as an unexamined panel. A complete blood count may support a search for infection, anaemia or systemic illness. Electrolytes, glucose, renal function, liver function, calcium and magnesium identify metabolic and organ-function disturbances that can impair cerebral function or alter drug handling.

Creatine kinase and inflammatory markers may be useful when examination suggests muscle injury, systemic inflammation or a severe behavioural state. Syphilis serology, medication levels and toxicology testing are selected when the history, risk context or unexplained course makes

them relevant. Urinalysis with culture can identify urinary pathology, but the result must still be related to symptoms and the whole clinical picture. An electrocardiogram forms part of medical assessment and is a safety requirement before some symptom-control choices. Chest radiography is included when respiratory or thoracic pathology is plausible.

INVESTIGATION DOMAIN	INCLUDED TESTS OR SAMPLES	CLINICAL QUESTION
Haematological and inflammatory	Complete blood count, inflammatory markers	Is there evidence of systemic illness, infection or tissue stress?
Metabolic and organ function	Electrolytes, glucose, renal and liver function, calcium, magnesium	Is a correctable physiological derangement present?
Muscle and medication-related	Creatine kinase, indicated medication levels, toxicology screen	Is there muscle injury, accumulation, intoxication or an exposure not captured by history?
Infective sampling	Urinalysis with culture, selected serology	Does the presentation warrant targeted investigation for infection?
Cardiorespiratory	Electrocardiogram, chest radiograph	Is there a medical contributor or a treatment-related cardiac constraint?

A normal result does not clear the patient of delirium, because delirium is diagnosed at the bedside and the relevant cause may be clinical, medication-related, evolving or outside the initial panel. Conversely, an abnormal result does not automatically establish causation. Rank abnormalities by their ability to explain the onset, severity and course, correct what is actionable, and observe whether the mental state changes as expected.

16.4 Conditional investigations: image or sample for a reason

Escalate beyond the core workup when history or examination points toward a neurological, respiratory or inflammatory cause. Lumbar puncture is indicated when central nervous system infection, an autoimmune disorder or subarachnoid haemorrhage is suspected. It is not a routine ritual for every confused patient. The decision follows assessment of whether the suspected condition is plausible and whether the procedure can be performed safely.

CT or MRI is selected when stroke is suspected or head trauma raises concern for haemorrhage. Imaging should answer a focused structural question. A psychiatric history must never lower the threshold for recognising focal findings or trauma, while a normal scan does not negate a clinically established delirium arising from systemic illness. Where immediate access differs, the available modality and urgency determine the practical sequence, but the clinical question remains the same.

Arterial blood gas analysis is appropriate when hypercarbia is a concern in chronic obstructive pulmonary disease. This is an example of hypothesis-driven testing: bedside respiratory assessment creates the question, and the test evaluates a physiological disturbance that routine

chemistry may not answer. Similar reasoning should govern every added investigation. Order a test because its result can revise the causal formulation or immediate management.

Detailed neuropsychological testing is usually deferred because impaired attention makes higher-order results difficult to interpret. This is not therapeutic neglect. Repeated bedside assessment is better suited to a fluctuating acute state; formal cognitive evaluation becomes meaningful after the patient can engage consistently. Standardised cognitive instruments and their scoring belong in the Psychotherapies, clinical assessment, MSE and nosology notes.

GOOD TO REMEMBER

Write the indication beside every conditional investigation. “CT because of focal findings after an acute change” is clinically stronger than “CT brain to rule out organicity,” which neither prioritises a diagnosis nor explains how the result will change care.

16.5 Definitive treatment and the non-pharmacological foundation

The decisive intervention is correction of the underlying medical illness, toxic exposure, medication effect or physiological derangement. This may require antimicrobial, respiratory, metabolic, surgical or medication-related treatment directed by the responsible medical team. Psychiatry should keep the causal plan visible in its note. Writing only “manage agitation” accidentally reframes an emergency as a behaviour problem.

Non-pharmacological support begins immediately and runs alongside definitive treatment. Reorient the patient calmly and repeatedly without arguing about every mistaken belief. Ensure access to appropriate sensory aids. Use familiar people and consistent explanations where possible. Support hydration and mobilisation according to medical safety, normalise the sleep-wake pattern, and provide reassurance. These measures reduce avoidable cognitive load and fear while preserving function.

The environment should make the correct behaviour easier. Keep communication brief, concrete and respectful. Introduce staff and explain procedures before touching the patient. Reduce unnecessary transfers, conflicting instructions and avoidable night-time disruption. Family can supply baseline information, familiar reassurance and observation, but should not be made solely responsible for containment. Nursing care remains active treatment: attention to comfort, elimination, nutrition, mobility, skin integrity and safe use of sensory aids prevents secondary burdens from perpetuating distress.

MNEMONIC

CAUSE first

Correct the precipitant; **A**id reorientation; **U**se sensory aids; **S**upport sleep, hydration and reassurance; **E**ncourage safe mobilisation. The mnemonic describes active treatment, not a prevention bundle.

Physical restraints should be avoided as far as possible. They can worsen psychotic symptoms and are associated with deep vein thrombosis, pulmonary embolism and cardiac dysrhythmias in

delirious patients. Continuous observation by staff or family may reduce the need for restraint. If immediate containment becomes unavoidable under local policy, the team should treat it as a temporary safety intervention, continue cause-directed care and reassess the need repeatedly.

16.6 The pharmacological threshold: rescue, not routine treatment

There is **no U.S. FDA-approved medicine for delirium or its behavioural disturbance**. Antipsychotic use is off-label, and studies in general medical and intensive-care settings have not shown reliable improvement in the outcomes that matter most, including severity, resolution, mortality or delirium duration. This evidence should stop routine prescribing, but it does not remove the need for proportionate rescue treatment when behaviour creates immediate danger or intolerable distress.

The indication is narrow: use a low-dose antipsychotic only when severe agitation or psychotic symptoms remain refractory to non-pharmacological care and create a safety risk. Choose the lowest effective dose, prescribe against a clearly documented target, and continue only while benefit outweighs risk. Taper and discontinue after symptom resolution. A calmer patient is not proof that the cause has been treated; sedation can make examination and fluctuation harder to observe.

Within the acute **LOCUS**, the **FOCUS** is a specific behaviour that prevents safe assessment or essential care, threatens the patient or others, or causes severe distress. The **MODUS** remains combined: staff presence, de-escalation, environmental support and cause-directed treatment continue while medicine is used. The prescription should state the delirium indication so it is not mistaken for treatment of a primary psychotic disorder during transfer or discharge.

TRAP **Do not write "haloperidol is the treatment of delirium."** It may control dangerous agitation or psychotic symptoms; it does not treat the precipitating illness, and routine antipsychotic use has not shown dependable delirium outcome benefit.

Antipsychotic receptor pharmacology and broader drug selection belong in the Schizophrenia: management, antipsychotics and adverse effects notes. Here the comparison is deliberately limited to delirium-scoped use, low dosing, route, tolerability and cardiac or neurological constraints.

16.7 Delirium-scoped antipsychotic dosing

Doses used for symptomatic control in delirium are low and indication-specific. Do not import a schizophrenia regimen. Oral treatment is preferable when the patient can take it safely and the clinical situation permits; intramuscular administration is an option when oral administration is not feasible and urgent behavioural control is necessary. Route choice also depends on monitoring capacity and the need to avoid coercion.

AGENT AND DELIRIUM INDICATION	DELIRIUM-SCOPED SCHEDULE	PRACTICAL CAUTION
Haloperidol, oral, for severe agitation or psychotic symptoms in delirium	**0.5-1 mg twice daily, with additional doses every 4 hours as needed**	Peak effect occurs at 4-6 hours; extrapyramidal effects are more likely above 3 mg.
Haloperidol, intramuscular, for the same delirium indication when oral treatment is not feasible	0.5-1 mg intramuscularly; observe for 30-60 minutes and repeat if necessary	Peak effect occurs at 20-40 minutes.
Quetiapine, oral, for severe agitation or psychotic symptoms in delirium in delirium	**12.5-50 mg twice daily; if tolerated, increase every 12 hours up to 200 mg daily**	Sedation and postural hypotension are common adverse effects.
Risperidone, oral, for severe agitation or psychotic symptoms in delirium in delirium	0.5 mg twice daily, with additional doses every 4 hours as needed; usual maximum 4 mg per day	Response may be poorer in older patients.
Olanzapine, oral, for severe agitation or psychotic symptoms in delirium in delirium	2.5-5 mg once daily, up to 20 mg per day	Sedation is the most commonly reported effect and may be useful when the presentation is markedly agitated.

For an older adult with delirium requiring behavioural control, in delirium, a conservative intravenous haloperidol schedule is 0.25-0.5 mg every 8 hours scheduled, with 0.25-0.5 mg every 6 hours as needed. If oral dosing is used, the corresponding schedule in delirium is 0.5-1 mg every 8 hours scheduled, with 0.5-1 mg every 6 hours as needed. If calm is achieved within 12-24 hours, continue for 2-3 days and taper over the next 2-3 days. This is a short symptom-control plan, not maintenance psychopharmacology.

Agent choice is secondary to indication discipline. Avoid combining agents merely because the first dose has not produced immediate tranquillity. Reassess the target, timing, route, accumulating adverse effects and the possibility that pain, hypoxia, fear or another untreated driver is sustaining the behaviour. Current local prescribing guidance and monitoring requirements should be checked before use.

16.8 QTc safeguards, exceptions and treatments to avoid

Obtain a baseline electrocardiogram before haloperidol. **Intravenous haloperidol QTc thresholds** are **450 ms or greater in a male and 470 ms or greater in a female**; at or beyond these values, intravenous haloperidol should not be used. Discontinue haloperidol if the QTc becomes prolonged by **more than 25% above baseline**. Interpret the tracing together with concurrent cardiac risk factors and other medicines rather than treating the number in isolation.

Antipsychotic sensitivity in Parkinson's disease and Lewy body dementia changes the plan: avoid haloperidol and antipsychotics generally. If a drug is genuinely needed here, a benzodiazepine such as lorazepam or diazepam is the usual alternative, used cautiously at the lowest effective dose because benzodiazepines can themselves worsen delirium; where an antipsychotic cannot be avoided, a low-dose agent with minimal dopamine blockade such as

quetiapine or clozapine is preferred under specialist guidance. Outside these neurological exceptions and sedative-hypnotic withdrawal, benzodiazepines should be avoided because they can exacerbate delirium, prolong it and increase falls risk, particularly in older adults. Withdrawal delirium, including its assessment and benzodiazepine regimens, belongs in the Substance use disorders notes.

Dexmedetomidine is setting-specific. Limited and mixed evidence suggests potential benefit in intensive-care populations, but cardiovascular adverse effects, particularly hypotension, confine its use to monitored intensive-care practice. Evidence should not be generalised to the ordinary ward, and no dose is supplied here because the permitted evidence does not provide one.

Cholinesterase inhibitors are not recommended for delirium. Adding rivastigmine to usual intensive-care treatment did not shorten delirium and was associated with a more severe presentation, longer intensive-care stay and higher mortality than placebo. Donepezil, rivastigmine and galantamine, used in Alzheimer disease, have not been adequately studied for delirium and should not be repurposed for it.

> Find and treat the precipitant, institute supportive care first, and reserve low-dose symptom medication for refractory severe distress or safety risk.

The safe endpoint is not simply a quiet patient. It is a patient whose cause has been identified as far as possible, whose physiology and comfort are being corrected, whose environment supports cognition, and whose temporary behavioural medication has a documented indication, monitoring plan and stopping decision. Failure to improve should reopen the causal formulation rather than automatically escalate sedation.

17 · Delirium prevention, outcomes, and substance-withdrawal delirium

Delirium prevention is active treatment of vulnerability before the syndrome appears. The marks lie in recognising that a substantial burden is avoidable, translating risk into a coordinated non-pharmacological bundle, and understanding that an episode is not necessarily transient in its consequences. Withdrawal delirium then supplies the crucial exception to ordinary symptomatic practice: the underlying withdrawal physiology must be treated, not merely the visible agitation.

> **IN INDIA**
>
> Give parenteral thiamine before or together with glucose in a malnourished or alcohol-dependent patient, without delaying treatment of hypoglycaemia, and treat alcohol-withdrawal delirium with a benzodiazepine rather than antipsychotic monotherapy.

This section names, but does not restate, delirium criteria and subtypes. It also uses the predisposing-versus-precipitating model only to explain why prevention is targeted; that model and delirium pathophysiology are taught in the dedicated aetiology section. Investigation and general management are addressed separately. Here the organising sequence is prevention, outcome surveillance, recognition of withdrawal, and safe exception management.

30% to 40%	OVER 200	35% to 40%	SIX-FOLD
CASES POTENTIALLY PREVENTABLE	HOSPITALS USING HELP	OLDER-PEOPLE MORTALITY AT ONE YEAR	MORTALITY ODDS WITH NEUROLEPTICS IN DT

17.1 Prevention begins with the vulnerability model

Multicomponent non-pharmacological prevention is the central approach. Delirium rarely reflects a single ward event that can be neutralised by a single tablet. Prevention instead reduces several modifiable pressures at the same time while protecting cognition, sleep, mobility, hydration, nutrition and sensory access. This is the practical consequence of the risk-factor model: a vulnerable brain may cross the threshold after several individually modest insults, so apparently ordinary elements of care become biologically important when they accumulate.

Among hospital cases, delirium may be preventable in **up to 30% to 40%**. This estimate makes prevention a patient-safety task rather than an optional refinement. The relevant question is not merely, “Does the patient have delirium now?” It is also, “Which foreseeable contributors can the team reduce before cognitive failure appears?” Screening identifies the population in whom prevention should be intensified, but the diagnostic criteria and screening instruments themselves belong in their dedicated delirium sections.

Prevention must be embedded in routine care. Risk is reviewed when the clinical context changes, and the bundle is adjusted to what the patient can safely do. Mobility, for example, is promoted within surgical and medical constraints; hydration and nutrition are supported in a manner compatible with the patient’s treatment; sensory aids help only if they are available and usable. This is coordinated individualisation, not a rigid checklist detached from the illness.

RULE **Prevent delirium by reducing cumulative avoidable burden, not by searching for a universal prophylactic drug.** The bundle acts across several modifiable domains because risk is usually produced by interacting vulnerabilities and insults.

17.2 HELP: the ward-based prevention bundle

The most widely used strategy is the **Hospital Elder Life Program (HELP)**. Its design is clinically coherent: maintain meaningful cognitive engagement, keep the patient moving, reduce unnecessary psychoactive exposure, protect sleep without relying on medication, maintain oral intake, and correct the avoidable isolation produced by impaired hearing or vision. Each domain removes a source of cognitive or physiological stress; together they create a ward environment less likely to push a vulnerable patient into delirium.

HELP DOMAIN	WARD ACTION	WHY IT MATTERS
Cognition	Reorientation and therapeutic activities	Preserves contact with place, people and purposeful activity
Mobility	Early mobilisation with exercise or walking	Prevents avoidable immobility from becoming part of the risk burden
Medicines	Reduce psychoactive-drug use and doses	Limits avoidable drug-related cognitive disturbance
Sleep	Promote sleep through non-pharmacological measures	Protects the sleep-wake cycle without adding sedative burden
Intake	Maintain oral hydration and adequate nutrition	Supports physiological reserve during illness
Senses	Provide appropriate hearing and vision aids	Improves access to orientation, communication and reassurance

MNEMONIC

Mind, Move, Medicines, Midnight, Meals, Missing senses

The alliterative spine recalls the HELP domains: cognitive reorientation and activity, early mobility, reduced psychoactive exposure, non-drug sleep promotion, hydration and nutrition, and appropriate hearing and vision aids.

HELP has been implemented in more than 200 hospitals worldwide and has shown effectiveness and cost-effectiveness in reducing delirium and improving outcomes for hospitalised older people. Scale matters because a prevention bundle succeeds only when it survives ordinary ward practice. Its strength is not a proprietary object or a single specialist encounter; it is a repeatable pattern of nursing, medical, rehabilitative and family-supported care.

17.3 Adapt the bundle to intensive care and surgery

In intensive care, prevention is organised through the **ABCDEF bundle**: assess and manage pain; coordinate spontaneous awakening and breathing trials; choose sedation and analgesia deliberately; assess and manage delirium; promote early mobility and exercise; and engage and empower family. The bundle is associated with a decreased risk of delirium. Its logic is continuity across disciplines: ventilation, sedation, pain, cognition, movement and family involvement cannot be managed as unrelated silos.

The perioperative lesson is similarly practical. In older adults undergoing hip-fracture repair, lighter sedation produced an **approximately 50% reduction in postoperative delirium risk**. Minimising benzodiazepines and selected opioids, particularly meperidine, can further reduce postoperative incidence. The point is not to deny necessary analgesia or anaesthesia. It is to avoid greater sedative exposure than the clinical task requires and to review the postoperative medication plan as part of delirium prevention.

Implementation should make omissions visible. The team can document whether the patient received the indicated orientation, mobility, sleep, intake and sensory measures, and why any element could not be delivered. When delirium nevertheless develops, this record helps distinguish an unavoidable limitation from a preventable care gap and allows the bundle to be revised. It also prevents the label "non-pharmacological" from being mistaken for passive care. These interventions require assessment, coordination, repetition and accountability.

- Place prevention in the care plan before cognitive change appears, rather than adding it after agitation begins.
- Assign each bundle element to a team member, because an unowned intervention is easily omitted.
- Reconcile psychoactive exposure across anaesthesia, intensive care, ward prescriptions and as-needed medication.
- Preserve mobilisation, sleep, oral intake and sensory access within the constraints of the medical condition.
- Use family engagement to support orientation and baseline knowledge, while retaining professional responsibility for care.

GOOD TO REMEMBER

The prevention prescription is written in verbs: orient, mobilise, review medicines, protect sleep, support intake and restore sensory access. "Observe for delirium" detects failure; it does not itself prevent it.

17.4 Why routine pharmacological prophylaxis is not the answer

Routine pharmacological prophylaxis has no clear evidence base. Evidence is limited for routine haloperidol or the second-generation antipsychotics olanzapine and risperidone, and there is no clear support for cholinesterase inhibitors or melatonin in non-intensive-care settings. Medication minimisation is itself an effective preventive strategy. Prescribing a psychoactive drug to compensate for preventable ward-level risk therefore reverses the preferred hierarchy.

A frequently examined signal requires careful interpretation. In an acute-care study published in 2014, ramelteon given for delirium prevention at 8 mg at night was associated with 3% delirium compared with 32% in the placebo group. Four of five later trials of melatonin analogues reported benefit, but heterogeneity limited broad recommendations. This is a prevention-outcome finding, not a delirium-treatment dose and not authority for routine use in every setting.

The apparent tension is resolved by evidence hierarchy and scope. A promising individual trial can justify further study and informed local discussion, while the total evidence can remain insufficient for universal prophylaxis. Candidates should state the signal and its limitation without turning uncertainty into nihilism: the finding is clinically interesting, but the standard prevention answer remains a multicomponent non-pharmacological bundle plus reduction of deliriogenic medication.

PITFALL

Do not convert the ramelteon study into routine prescribing, and do not mislabel 8 mg at night as treatment for established delirium. The cited regimen concerns prevention in acute-care patients.

17.5 Outcomes: delirium is not a transient inconvenience

Clinical recovery from the acute behavioural state does not erase prognosis. In older people with delirium, the **one-year mortality is 35% to 40%**. The excess mortality risk persists at 12, 24 and 36 months, with a risk ratio of at least 2 at each time point; increased cognitive and functional impairment also remains evident at 24 months. These associations reflect a population with serious illness and vulnerable brains, but they still make delirium an important prognostic marker.

Persistence must be described by the outcome actually measured. A course account in the comprehensive psychiatry textbook reports residual delirium symptoms in 20% of older adults for up to 6 months; another overview in that textbook reports as many as 15% still symptomatic at 6 months. By contrast, the physical-illness prescribing guideline reports persistent cognitive impairment after delirium in up to 60% of individuals. Residual delirium and residual cognitive impairment are not interchangeable constructs, and their percentages must not be averaged or presented as competing estimates of the same endpoint.

OUTCOME CONSTRUCT	REPORTED FINDING	CORRECT INTERPRETATION
Persistent delirium symptoms	20% of older adults with residual symptoms for up to 6 months; another account gives as many as 15% symptomatic at 6 months	The syndrome or some of its symptoms remain present
Persistent cognitive impairment	Up to 60% following delirium	Cognitive deficit persists after the episode; this is broader than residual delirium
Mortality in older people	35% to 40% at one year	Delirium identifies a group requiring serious prognostic attention
Longer-term mortality association	Risk ratio at least 2 at 12, 24 and 36 months	The mortality association is sustained beyond discharge

TRAP **Do not collapse persistent delirium into persistent cognitive impairment.** The 15% and 20% reports concern residual delirium symptoms, whereas up to 60% concerns cognitive impairment after delirium. The denominators and measured outcomes differ.

17.6 Cognitive follow-up: delirium as a cerebral stress test

In vulnerable older adults, delirium may leave permanent cognitive impairment and increase incident dementia; affected patients have been reported to be three times more likely to develop dementia. Patients with Alzheimer disease may decline more rapidly after an episode. The relationship is bidirectional: pre-existing cognitive vulnerability increases susceptibility to

delirium, while delirium may accelerate subsequent decline. This is why an apparent return to behavioural calm should not end cognitive and functional observation.

The useful clinical metaphor is a **brain stress test**. Serious illness may expose cognitive deficits that were present but not obvious in ordinary life. The episode should therefore trigger careful reconstruction of premorbid function, review of recovery with an informant, and attention to whether cognition and daily functioning return toward baseline. It should not trigger a dementia diagnosis during an unstable confusional state. The dementia syndrome and its full workup belong in the Dementias and neurocognitive disorders notes.

Outcome assessment should separate several questions: Has the delirium itself resolved? Is cognition still impaired? Has mobility or self-care fallen below baseline? Is the family observing a new trajectory? Does the patient now require a different level of support? This separation prevents a quiet, dependent patient from being labelled “recovered” merely because disruptive behaviour has ceased. It also avoids treating every post-illness cognitive complaint as proof of a progressive neurocognitive disorder.

Discharge communication should preserve diagnostic uncertainty and the need for review. State the pre-illness baseline as far as known, the degree of recovery achieved, and the domains that remain impaired. The long-term outcome data justify follow-up, but they do not predict the fate of an individual patient with certainty.

Prognostic communication should be serious without becoming deterministic. Families can be told that lingering confusion and cognitive or functional decline are recognised outcomes, that recovery may be incomplete, and that a changed trajectory deserves reassessment. The clinician should also avoid attributing every later difficulty directly to delirium, because the underlying illness and premorbid vulnerability remain part of the causal picture. The episode is an event to recover from as well as a marker that future stress may expose limited cerebral reserve.

17.7 Recognising alcohol-withdrawal delirium within the differential

Delirium tremens is the severe agitated delirium of alcohol withdrawal. It occurs in around **3% to 5% of people admitted to hospital for alcohol withdrawal** and typically develops around 72 hours after the last drink. A careful alcohol history must therefore establish the last reliable intake, subsequent access to alcohol, prior severe withdrawal and the chronology of autonomic and behavioural change. A patient may be admitted for another illness and develop withdrawal only after alcohol access stops.

The broader withdrawal syndrome provides context without making delirium the expected outcome. Fewer than 10% of people who develop alcohol withdrawal ever develop dramatic manifestations such as severe autonomic hyperactivity, marked tremor or withdrawal delirium, and tonic-clonic seizures occur in fewer than 3%. Delirium tremens lies at the severe end of the alcohol-withdrawal spectrum and may emerge later than the initial withdrawal symptoms; its detailed symptom time course belongs in the Substance use disorders notes.

Previous seizures or delirium, low potassium, low magnesium, thiamine deficiency, systemic disease and undertreated withdrawal predispose to delirium tremens. These features increase vigilance; they do not replace bedside diagnosis or justify waiting for the full syndrome before treating escalating withdrawal.

The full withdrawal timeline, the content and scoring of the **Clinical Institute Withdrawal Assessment for Alcohol, revised (CIWA-Ar)**, and complete benzodiazepine regimens belong in the Substance use disorders notes. Within the delirium differential, the essential task is to recognise the withdrawal context because it changes the pharmacological principle.

17.8 Withdrawal delirium: LOCUS, FOCUS and MODUS

LOCUS: delirium tremens is a medical emergency requiring general-hospital care. NICE guideline CG100, updated in 2017, recommends nursing the patient in a general hospital and using lorazepam orally or intravenously. The setting must support medical assessment, repeated observation, correction of contributing illness and escalation when withdrawal remains uncontrolled.

FOCUS: in the acute phase, the target is the withdrawal state and its associated agitation, autonomic activation and risk, not tranquillity alone. As control is achieved, focus moves to preventing recurrence, correcting contributors and reassessing whether another cause of delirium is also present. A familiar low-dose rapid-tranquillisation protocol, particularly one that includes haloperidol, is explicitly inappropriate for delirium tremens because it is designed around a different clinical problem.

MODUS: **benzodiazepines are the medications of choice** for alcohol or sedative-withdrawal delirium. Delirium tremens may require larger benzodiazepine doses than ordinary rapid-tranquillisation practice. In reduced hepatic metabolic capacity, lorazepam or oxazepam may be preferred to chlordiazepoxide because they have shorter half-lives and do not depend on hepatic oxidative metabolism. Route and intensity are matched to urgency, clinical response and the capacity for monitoring.

Antipsychotics are not equivalent treatment. They do not correct the underlying GABA, glutamate and noradrenergic disturbance of withdrawal delirium. More importantly, a meta-analysis comparing sedative-hypnotics with antipsychotics in delirium tremens found an **approximately six-fold increase in the odds of mortality with neuroleptics**. This safety signal makes antipsychotic monotherapy a dangerous category error, not merely a weaker option.

The withdrawal-specific scoring and dosing, including any CIWA-Ar based titration, belong in the Substance use disorders notes and current local guidance; CIWA-Ar depends on patient participation and is in any case unreliable once established delirium prevents valid self-report. The delirious patient also requires the ordinary medical work of examining for concurrent illness, reviewing electrolytes and nutrition, maintaining supportive care and reassessing the diagnosis when the course is atypical. Withdrawal may be the central cause without being the only cause.

MODUS also includes non-drug care. Psychological support is brief, concrete and repeated: reduce stimulation, explain what is happening, reorient without confrontation and use calm human presence alongside medical treatment. Procedural care comprises the monitoring and access required to deliver treatment safely and respond to deterioration. Social input includes collateral history about alcohol and sedative use, involvement of a suitable caregiver, and continuity planning after the emergency settles. None of these modes replaces benzodiazepine treatment of the withdrawal state; they make that treatment safer and the surrounding care more complete.

In alcohol or sedative-withdrawal delirium, treat the withdrawal with benzodiazepines; antipsychotic monotherapy is not equivalent management and may be dangerous.

18 · Korsakoff syndrome: pathology, mammillary body and thalamic involvement

Korsakoff syndrome is the model diencephalic amnesia. The marks do not come from writing "alcohol causes memory loss." They come from linking thiamine deficiency to bilateral mammillary and thalamic pathology, placing those lesions within a distributed memory circuit, and then explaining why new learning fails while immediate recall and procedural learning can remain available. A strong answer therefore moves in a chain: antecedent deficiency, pathological structures, disrupted circuit, clinical memory profile, imaging correlate and prognosis.

IN INDIA

Give parenteral thiamine before or together with glucose in any at-risk patient, without delaying treatment of hypoglycaemia; the cost of missing it, a persistent and often irreversible amnestic syndrome, is out of all proportion to the simplicity of giving it.

The central anatomical distinction is easy to blur. The dorsomedial thalamic nuclei and mammillary bodies are sites of haemorrhagic pathology in the acute antecedent encephalopathy, while the anterior thalamic nuclei receive mammillary input and participate in an explicit-memory network with medial temporal and cingulate structures. Both thalamic regions belong in the answer, but they do not have identical circuit roles. The accompanying account keeps pathology, connectivity and clinical expression at their proper levels.

0.4% TO 2%	2%	25%
BOTH SYNDROMES, COUNTRY-SPECIFIC	WERNICKE, GENERAL POPULATION	KORSAKOFF RECOVERY

18.1 Definition, antecedent and diagnostic language

In 1889, **Sergei Korsakoff** described a syndrome combining polyneuritis, anterograde amnesia and confabulations in people with chronic alcohol use. The description did not separate the acute illness from the enduring memory disorder; Carl Wernicke had described the acute encephalopathy earlier. This history explains why the paired term Wernicke-Korsakoff syndrome remains common, but it must not make the candidate treat acute encephalopathy and chronic amnesia as interchangeable states.

Korsakoff syndrome usually develops after Wernicke encephalopathy, although some patients develop the chronic syndrome without a recognised initial acute stage. Absence of a documented acute episode therefore does not abolish the pathological link to thiamine deficiency. It may mean that the antecedent was not clinically recognised, but the permitted evidence does not quantify how often this occurs. The acute thiamine-deficiency emergency and its full treatment protocol are covered in the Endocrine and metabolic psychiatry notes.

In the setting of severe alcohol use disorder, **DSM-5** refers to the condition as alcohol-induced major neurocognitive disorder, amnestic-confabulatory type. This label describes the alcohol-related setting and the dominant cognitive phenotype. It should not be used to erase non-alcoholic thiamine deficiency, nor should it turn the bedside formulation into a mere coding exercise. The defining clinical problem remains a persistent amnestic syndrome arising from damage to a memory network.

RULE **Write the sequence, not a fused label.** Wernicke encephalopathy is the usual acute antecedent; Korsakoff syndrome is the enduring amnestic state, and the latter can appear without a recognised acute presentation.

18.2 Thiamine deficiency: the metabolic cause

Thiamine pyrophosphate is the active form of thiamine and a cofactor for enzymes central to glucose metabolism and neurotransmitter production. Its lack compromises the energy metabolism on which vulnerable diencephalic neurons depend, which is why deficiency itself, and not poor nutrition as a label, is the operative cause of the amnestic pathology.

Alcohol is the common context but not an obligatory one. Severe thiamine deficiency also arises from general malnutrition and from several non-alcoholic medical settings, which the Endocrine and metabolic psychiatry notes cover alongside the acute thiamine emergency. The clinical implication is diagnostic discipline: a patient need not fit the stereotype of long-standing alcohol misuse for thiamine-deficiency brain injury to be plausible, and chronic alcohol use does not by itself prove that an amnestic presentation is Korsakoff syndrome. The history supplies the deficiency context; the memory phenotype, neurological course, imaging and exclusion of competing causes build the localisation.

18.3 The pathological core: mammillary bodies and thalamus

The acute pathological substrate is **bilateral haemorrhage in the mammillary bodies and dorsomedial thalamic nuclei** in Wernicke encephalopathy secondary to thiamine deficiency. When chronic alcohol-related thiamine deficiency produces permanent injury to the mammillary bodies and thalamus, the resulting Korsakoff syndrome may be irreversible. The enduring memory disorder is therefore not adequately explained as intoxication, poor effort or a vague diffuse cognitive decline.

The anatomical field is wider than those two headline structures. Thiamine-deficiency lesions also include the fornix and cerebellum, and both the medial dorsal and anterior thalamic nuclei have been associated with the deficiency state. The mammillary complex and anterior thalamic nuclei are particularly important for explicit memory. The fornix matters because it carries

hippocampal fibres toward the mammillary bodies, while the anterior thalamic nucleus carries the mammillary signal onward toward cingulate components of the memory circuit.

Two distinctions protect the answer from overstatement. First, "thalamic involvement" is not a single undifferentiated dot. The dorsomedial nucleus is named in the haemorrhagic pathology; the anterior nucleus is named as a thiamine-deficiency lesion site and as a connected relay for explicit memory. Second, the permitted evidence does not adjudicate which thalamic nucleus is the unique critical lesion for diencephalic amnesia. The safest high-level formulation is network injury involving mammillary and thalamic structures, with disruption of their connections to medial temporal and cingulate regions.

TRAP **Do not choose between "mammillary bodies" and "thalamus," or between dorsomedial and anterior thalamus.** Thiamine-deficiency pathology includes mammillary, medial dorsal and anterior thalamic structures, while the anterior nucleus has the explicit mammillothalamic circuit role.

18.4 The mammillothalamic memory circuit

The circuit of Papez provides the anatomical spine for understanding the amnesia. Fibres arise in the mammillary body, travel through the mammillothalamic tract and reach the anterior thalamic nucleus. The anterior nucleus projects to the cingulate cortex and cingulum, which connect onward with entorhinal and subicular cortices on the medial temporal surface. From there, the pathway enters the hippocampal formation and eventually returns through the fornix to the mammillary body.

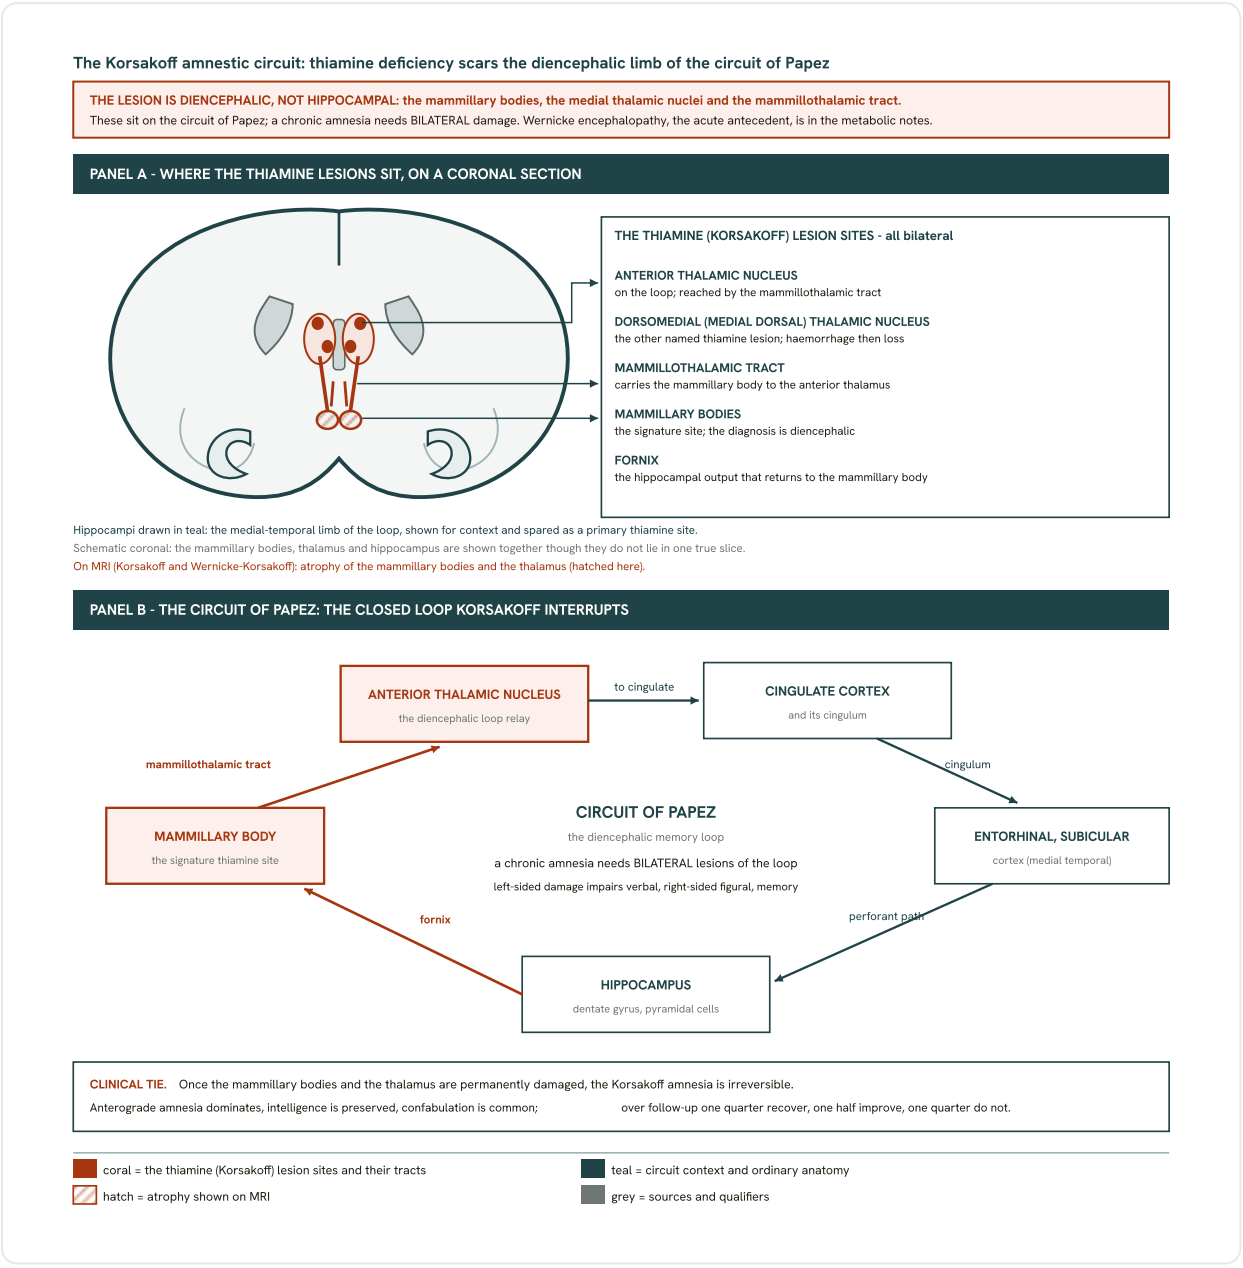


Fig 34.6 **The Korsakoff amnestic circuit.** Thiamine deficiency scars the diencephalic midline, the mammillary bodies, the medial (dorsomedial and anterior) thalamic nuclei and the mammillothalamic tract, which form the diencephalic limb of the circuit of Papez. A chronic amnesia needs bilateral lesions of the loop, and once the mammillary bodies and thalamus are permanently damaged it is irreversible. The hippocampus is the medial-temporal limb, shown for context and spared as a primary thiamine site. (Kaplan and Sadock CTP 2024, and Tasman Psychiatry.)

The accompanying figure should be read as a loop rather than a list of unrelated nuclei. The entorhinal cortex projects through the perforant pathway to the dentate gyrus; fibres then reach the pyramidal cells of the hippocampus proper, and hippocampal fibres enter the fornix before terminating in the mammillary body. The loop links diencephalic structures with the medial temporal memory system and cingulate cortex. Damage at a relay or connecting tract can therefore disturb explicit memory even when the entire circuit is not destroyed.

NODE OR TRACT	SUPPORTED CONNECTION	WHY IT BELONGS IN THE ANSWER
Mammillary body	Origin of the mammillothalamic tract; receives the returning fornix pathway	Major site of thiamine-deficiency pathology
Mammillothalamic tract	Mammillary body to anterior thalamic nucleus	Part of the anterior-thalamic pathway damaged in amnesia
Anterior thalamic nucleus	Projects to cingulate cortex and cingulum	Connected with explicit-memory impairment when injured
Entorhinal, subicular and hippocampal structures	Perforant pathway to dentate gyrus and hippocampus	Medial temporal portion of the loop
Fornix	Carries hippocampal fibres back toward the mammillary body	Also named among thiamine-deficiency lesion sites

MNEMONIC

MAMMILLARY - TRACT - ANTERIOR THALAMUS - CINGULATE - MEDIAL TEMPORAL - FORNIX

Follow the loop from **mammillary** body through the mammillothalamic **tract** to **anterior thalamus**, then **cingulate** and **medial temporal** structures, before the **fornix** carries the return pathway.

18.5 Network logic, bilaterality and material-specific laterality

Bilateral circuit injury is usually required for chronic amnesia involving the circuit of Papez, although exceptions occur. This principle fits the bilateral pathological emphasis in thiamine-deficiency disease. It also explains why the syndrome is better conceived as disconnection of a distributed memory network than as loss of a single isolated "memory centre." Mammillary injury, thalamic injury and tract interruption can converge on the same failure of new explicit learning because they interrupt different positions in the loop.

Bilaterality does not mean that hemispheric specialisation disappears. Left-sided hippocampal or thalamic damage produces predominantly verbal memory impairment, whereas right-sided damage produces predominantly figural memory impairment. These lateralised deficits describe the type of material most affected after unilateral or asymmetric damage. The chronic global amnestic syndrome, by contrast, generally points toward bilateral disruption. Thus, "left verbal, right figural" and "bilateral for chronic amnesia" answer different localisation questions.

The same evidence clarifies why immediate retention can survive. Short-term or working memory depends on frontoparietal mechanisms distinct from the hippocampal and thalamic pathways damaged in the amnestic syndrome. A patient may therefore hold a small amount of information briefly yet fail to establish a durable new memory from it. This apparent contradiction is diagnostically useful: intact immediate performance does not demonstrate intact learning across a delay.

This dissociation also prevents a common bedside misreading. A coherent conversation samples attention, language, general knowledge and immediate maintenance as well as memory, so conversational fluency can mask the amnestic deficit. The examiner must introduce material, allow the immediate moment to pass and then test whether it was retained. Failure of delayed new learning despite preserved immediate recall is more informative than repeatedly asking the patient whether memory feels poor. Likewise, improvement on a practised procedure without conscious recollection of the learning episode demonstrates preserved implicit learning rather than recovery of explicit episodic memory. The network account thus predicts a selective pattern, not uniform failure of every task containing the word "memory."

The network model also disciplines interpretation of imaging. Atrophy in a characteristic structure supports the formulation, but a scan is not a direct photograph of a cognitive operation. The clinician still has to show that the bedside pattern is amnestic, that the relevant history supports thiamine deficiency, and that alternative causes have been considered. Anatomy explains the pattern; it does not replace phenomenology.

18.6 The clinical memory profile

The dominant deficit is **anterograde amnesia**: a striking inability to form new memories despite broadly intact overall intelligence. The failure becomes clearest when information must be retained beyond immediate registration and later retrieved as an episode. The patient may participate coherently in the interview, repeat material in the moment and appear socially intact, yet be unable to retain what has just been discussed.

Several capacities are characteristically spared, which sharpens the contrast with the failure of new episodic learning:

- immediate registration and repetition within the moment;
- procedural and skill learning acquired without conscious recollection;
- broadly preserved general intelligence and language;
- a socially intact surface that can mask the severity of the deficit.

Retrograde amnesia is present but less severe, and it usually affects events shortly before the illness began. This temporal pattern should not be flattened into "both recent and remote memory are lost." The asymmetry matters: new learning is the striking impairment, while loss of previously formed memories is more limited. The exact boundary of lost autobiographical material must be established clinically rather than inferred from the label.

MEMORY DOMAIN	TYPICAL FINDING	INTERPRETIVE MEANING
New explicit learning	Strikingly impaired	Core anterograde amnesia
Memory for pre-illness events	Less severely impaired, especially near onset	Retrograde component is present but subordinate
Immediate recall	Preserved	Working-memory mechanisms are relatively distinct
Implicit or procedural learning	Preserved	Performance can improve without normal explicit recollection
Overall intelligence	Broadly intact	The syndrome is not synonymous with global intellectual failure

Confabulation is a classical accompanying feature and is commonly understood as an attempt to cover the memory deficit. Its presence may draw attention to the syndrome, but its phenomenological types and mechanisms are taught in the dedicated confabulation section. It should not displace direct examination of acquisition, retention and delayed retrieval.

GOOD TO REMEMBER

Separate registration from learning. Preserved immediate recall can coexist with profound failure to form new memories, and preserved procedural learning can coexist with poor conscious recollection of having practised. Test performance across a delay rather than concluding from an intact immediate response.

18.7 Imaging and diagnostic integration

Mammillary-body and thalamic atrophy is the typical imaging correlate in Wernicke-Korsakoff syndrome. Mammillary-body atrophy is particularly useful as an anatomical clue within the imaging differential of amnestic syndromes. The finding coheres with the pathological account: structures that link hippocampal output with anterior thalamic and cingulate components of the memory network have undergone enduring injury.

Imaging is supportive and discriminative, not self-sufficient. Neuroimaging helps distinguish several causes of amnestic syndrome, but the alternative structural and transient amnesias are covered in their dedicated section. For Korsakoff localisation, the correct inference is narrower: mammillary and thalamic atrophy supports a thiamine-deficiency diencephalic amnesia when the history and memory phenotype agree. Absence of a quoted sensitivity or specificity prevents a stronger diagnostic claim.

At the bedside, build the formulation in layers. Identify a disproportionate failure of new learning; verify that immediate recall, procedural learning and overall intellectual function are not being mistaken for normal episodic memory; look for the less severe retrograde component; and integrate nutritional or alcohol-related risk with imaging. The full ordered cognitive examination and standardised cognitive instruments belong in the Psychotherapies, clinical

assessment, MSE and nosology notes. The dementia syndrome and its broader workup belong in the Dementias and neurocognitive disorders notes.

PITFALL

Do not diagnose the syndrome from alcohol history plus forgetfulness. The localisation becomes persuasive when severe anterograde amnesia coexists with preserved immediate recall and implicit learning, and when mammillary or thalamic atrophy supports the expected diencephalic network injury.

18.8 Prevention, treatment boundary and prognosis

The decisive therapeutic principle lies before fixed amnesia is established. Vitamin B1, particularly by the intravenous route, can prevent cognitive impairment and Wernicke-Korsakoff syndrome in people at risk because of long-standing alcohol misuse or severe malnutrition. In an emergency or inpatient locus, the immediate focus is recognition and treatment of the acute deficiency state; in later care, the focus shifts to persistent memory impairment, safety and support. The exact acute repletion regimen belongs in the Endocrine and metabolic psychiatry notes.

The treatment boundary matters because high-dose intravenous thiamine is likely to reverse Wernicke encephalopathy partly or completely, whereas benefit in established chronic Korsakoff syndrome is not established. This does not make late care futile. It means that thiamine treatment of the acute state and rehabilitation of chronic memory disability are different targets. Drug treatment addresses the deficiency emergency; psychological, environmental and social modes must then be aligned with the enduring functional consequences, although the permitted evidence does not provide a validated rehabilitation programme.

Long-term outcome is heterogeneous. **About one-quarter recover**, one-half improve to some degree, and the remaining one-quarter show no recovery. A parallel source summarises recovery at about 25% and notes that the others retain varying degrees of memory impairment. A significant minority require permanent placement. These figures are best presented as an outcome distribution, not as a prediction for an individual patient.

The prognostic message follows directly from the pathology. Permanent mammillary-body and thalamic damage can make the syndrome irreversible. Prevention and early treatment of the antecedent deficiency therefore carry disproportionate value. Once chronic amnesia is present, care planning should not be based on an assumption of either inevitable full recovery or inevitable complete non-response. The observed range includes recovery, partial improvement and persistent impairment.

Epidemiological figures require equally careful scope. Reported rates for Wernicke encephalopathy and Korsakoff syndrome range from 0.4% in France to 2% in the United States, Scandinavia and Australia. A separate estimate places Wernicke encephalopathy prevalence in the general population at 2%. These are distinct estimates and should not be merged into a single

universal rate. The syndrome is usually seen in people over 40 years of age after 20 or more years of alcohol use, but non-alcoholic deficiency remains clinically important.

> Korsakoff syndrome is a predominantly anterograde amnesia caused by permanent disruption of a bilateral mammillary-thalamic memory network, with immediate recall and procedural learning relatively preserved.

19 · Confabulation: types and mechanisms

Confabulation is a disorder of remembering in context, not merely an inaccurate statement. The examiner is looking for a disciplined distinction between accounts elicited by a memory question and accounts produced spontaneously, followed by an explanation of why frontal monitoring failure changes their content, persistence and resistance to correction. The marks lie in describing the behaviour before naming a lesion: what prompted the account, what kind of narrative appeared, how the patient responded to contradiction, and whether the anatomical setting fits the proposed type.

The useful clinical stance is descriptive. An account may conflict with the verified situation, yet its surface falsity alone does not explain the cognitive process that produced it. The psychiatrist should establish the memory context, compare the account with collateral information, and observe belief revision without turning the interview into an argument. The core distinction comprises **two** patterns: **provoked confabulation** and **spontaneous or fantastic confabulation**. They overlap enough that no single sentence should classify the phenomenon; the pattern across prompting, narrative quality and correction is more informative.

PROMPT	CONTENT	CHECK	LOCUS
WHAT ELICITED IT?	PROSAIC OR FANTASTIC?	DOES BELIEF REVISE?	DOES ANATOMY FIT?

19.1 What the term means at the bedside

Confabulation in chronic amnesia becomes visible when a patient supplies an account that conflicts with the established circumstances. The account is clinically meaningful because it is offered as memory, not because it is dramatic or bizarre. A routine answer about recent activity can therefore be more informative than an elaborate story. The clinician first asks what memory demand preceded it, whether the answer is anchored to the current setting, and how readily it changes when accurate contextual information is supplied.

This approach prevents several conceptual errors. Confabulation should not be reduced to ordinary factual inaccuracy; the relevant question is how an erroneous recollection is generated and accepted. Implausibility is not required, because a simple, everyday account may be a classic provoked example. Plausibility likewise does not establish accuracy. Verification comes from records, staff observation and collateral history, while phenomenology comes from the patient's manner of production and response to correction.

The word is also not a complete diagnosis. It describes a memory-related phenomenon whose clinical value depends on the surrounding amnestic state and neuroanatomical context. The

pathology of Korsakoff syndrome belongs in the dedicated section on mammillary-body and thalamic involvement. Other causes of amnestic syndrome and the transient amnesias are discussed separately. The full ordered cognitive examination and standardised cognitive instruments belong in the Psychotherapies, clinical assessment, MSE and nosology notes; here the examination is limited to the operations needed to characterise confabulation.

RULE **Classify the production pattern, not the truth value alone.** Prompting, content, response to correction and lesion context together distinguish the clinically useful forms.

19.2 Provoked confabulation

Provoked confabulation is typically elicited by a question. The examiner asks about a recent event, and the patient provides an answer despite the account being inconsistent with what actually occurred. Its content is usually **prosaic, simple and straightforward**. Ordinary routines, familiar social activity or other plausible autobiographical material may be recruited into the answer. Because the narrative resembles everyday life, it can escape notice unless the interviewer has reliable knowledge of the patient's actual circumstances.

The second defining feature is **little resistance to correction**. When the discrepancy is pointed out calmly and the true context is supplied, the patient may pause, reconsider and accept that the initial account was probably wrong. The correction does not necessarily restore the missing memory. What changes is the patient's commitment to the proposed account. This distinction matters: revisability shows that the account is not being defended with the persistence typical of the spontaneous pattern.

Provoked confabulation can be understood as an error exposed by retrieval demand. A question requires an answer about a time and situation that the damaged memory system cannot adequately reconstruct. Familiar material may be selected because it is coherent and available, but its fit to the requested episode is poor. The resulting narrative may therefore be internally sensible while externally misplaced. The interviewer should not mistake coherence for accurate episodic placement.

At the bedside, use neutral questions before offering alternatives. Record the exact question, the gist of the response, the verified context and the reaction to correction. Repeated leading questions can manufacture apparent consistency by repeatedly supplying the same cues. A calm correction is more informative than confrontation because the observation sought is whether the patient can update the account when contextual evidence becomes available.

19.3 Spontaneous or fantastic confabulation

Spontaneous confabulations often emerge independently of prompting, although they may also appear in response to a question. Their defining contrast extends beyond speech beginning uninvited. The narratives often have a **fantastic or dream-like quality**, extending beyond the ordinary substitutions seen in provoked confabulation. A patient may describe an implausible plan as if it belongs to the current situation and may organise behaviour around that account.

Correction is characteristically resisted. When the clinician supplies the verified context, the patient may become irritated and continue to insist on the account. Resistance is therefore part of the phenomenology, not an interpersonal nuisance to be ignored. It suggests that the false account has passed whatever internal process would normally reject material that is temporally misplaced, contextually inappropriate or incompatible with current reality.

The classic clinical setting in the permitted evidence is a patient assessed after clipping of an anterior communicating artery aneurysm, in the context of basal forebrain damage. The value of this example is anatomical and phenomenological. The narrative is spontaneously relevant to action, fantastic in content and maintained despite direct contradiction. The clinician should learn the pattern rather than reproduce the anecdote as a substitute for explanation.

Spontaneous confabulation occurs in only a minority of amnestic patients. Amnesia alone is therefore insufficient as a mechanism. The additional clue is failure of frontal systems that normally evaluate when and from where a recollection arose, and whether it is appropriate to the present situation. The account is not merely retrieved inaccurately; it is insufficiently checked before it is accepted and acted upon.

PITFALL

Do not label every implausible statement in an amnestic patient as spontaneous confabulation. Establish whether it emerged as memory, whether it was produced without prompting, whether its quality was fantastic or dream-like, and whether it persisted when the verified context was supplied.

19.4 The comparison that earns marks

The highest-yield discrimination rests on **three** bedside axes: prompt dependence, narrative quality and response to correction. Localisation then tests whether the classification is coherent. None of these axes should be read in isolation. Spontaneous confabulation can occasionally follow a question, so an elicited response is not automatically provoked. Conversely, a plausible account is not necessarily accurate. The whole production pattern decides the label.

CLINICAL AXIS	PROVOKED CONFABULATION	SPONTANEOUS OR FANTASTIC CONFABULATION
How it appears	Usually elicited by a memory question	Often volunteered, but may follow a question
Narrative quality	Prosaic, simple and straightforward	Fantastic and sometimes dream-like
Correction	Little resistance; the account is readily reconsidered	Resistance and irritation may occur; the account is maintained
Lesion association	May accompany lesions anywhere in the circuit of Papez	Generally associated with bilateral basal forebrain lesions
Interpretive emphasis	Retrieval under demand with weak episodic fit	Retrieval plus failure to monitor context and reject the account

MNEMONIC

PROMPT - PLOT - PUSHBACK

Prompt: was the account elicited or volunteered? **Plot:** was it ordinary or fantastic? **Pushback:** did accurate contextual information lead to revision or resistance? Use lesion localisation as the coherence check after describing these features.

The table also shows why "confabulation equals lying" is not an adequate examination answer. Truth value is only the starting observation. The diagnostic work is to analyse the memory operation and the checking failure. Equally, "confabulation equals gap filling" is too vague unless the candidate states which form, what elicited it and how the patient responded when reality was clarified.

19.5 Lesion localisation and the amnesia substrate

Provoked confabulation may occur with chronic amnesia caused by lesions anywhere in the **circuit of Papez**. It is most commonly encountered in Korsakoff syndrome secondary to thiamine deficiency. This broad association fits the clinical picture: disruption within an established memory circuit can leave the patient unable to retrieve the requested episode accurately, yet does not by itself imply the more severe monitoring disturbance that sustains fantastic accounts.

By contrast, spontaneous confabulation generally occurs in chronic amnesia associated with **bilateral basal forebrain lesions**. The laterality and regional restriction are important. They explain why spontaneous confabulation is not simply the more colourful end of ordinary amnesia. It reflects an amnestic substrate plus damage in a region closely related to frontal systems that evaluate relevance, temporal placement and behavioural meaning.

These anatomical statements should be held together, not memorised as disconnected lists. A distributed memory-circuit lesion can support provoked errors. The spontaneous form points more specifically toward bilateral basal forebrain involvement and its relationship with ventromedial frontal monitoring. Thus localisation follows phenomenology: first decide what kind of account was observed, then ask whether the lesion pattern is expected. Working backwards from a scan to a behavioural label risks circular reasoning.

Confabulations are more frequent in **diencephalic amnesia**, exemplified by Korsakoff syndrome, than in **hippocampal amnesia**. This relative-frequency statement reinforces that equally severe-looking memory failure need not produce the same confabulatory tendency. It does not justify diagnosing a lesion from confabulation alone. Rather, it tells the candidate that the organisation of the damaged memory system matters, not merely the presence of forgetting.

19.6 Mechanism: source and temporal-context failure

The most economical mechanism begins with **source memory**: the ability to identify where information came from and the circumstances to which it belongs. A recollection may contain familiar people, places or actions yet be assigned to the wrong episode. If its origin is not

recovered, familiarity can be mistaken for evidence that the event occurred in the requested context.

Frontal damage impairs memory for temporal context and source, and this deficit is relevant to confabulation. Temporal-context memory answers “when did this belong?” Source memory answers “where did this information arise?” A patient may retain fragments of autobiographical knowledge while losing the tags that separate past routine, imagined plan and current event. The error is therefore not necessarily absence of content. It may be failure to bind available content to its proper context.

This model separates availability from attribution. Material may feel familiar because it belongs somewhere in the person’s experience, yet familiarity does not specify its source or time. The memory system must still decide whether the material was perceived, planned, imagined or recalled from another occasion, and whether it answers the question now being asked. Failure at that stage creates an account whose components can be individually ordinary while their assembly is wrong for the present episode. A further monitoring failure allows the assembled account to survive obvious conflict with the current setting.

Correction tests the final part of this sequence. Acceptance suggests that external context can still overrule the initial reconstruction. Resistance suggests that the internally generated account continues to govern judgment despite contrary information. The response therefore gives the examiner a behavioural window into monitoring: it does not directly measure the missing memory, but it shows whether the patient can use supplied evidence to revise what has just been claimed.

The next operation is evaluation. Retrieved material must be checked against the present environment, recent experience and other available evidence. When this monitoring works, an initially plausible candidate memory can be rejected or held with uncertainty. When it fails, the candidate account may be accepted as current fact. Provoked confabulation can be conceptualised as weak reconstruction under demand with enough checking capacity remaining for later revision. Spontaneous confabulation adds a more consequential failure of monitoring, so contextually inappropriate material survives contradiction.

Spontaneous confabulation depends on **ventromedial frontal damage**. This localisation is consistent with, rather than opposed to, the bilateral basal forebrain association. The ventromedial frontal and basal forebrain region provides the anatomical bridge between memory failure and defective evaluation. The mechanism therefore has paired levels in the answer: an amnestic disturbance supplies poorly contextualised material, while frontal monitoring failure permits that material to be believed, maintained and potentially translated into action.

19.7 Bedside elicitation and interpretation

A useful assessment samples autobiographical material whose correct context can be independently established. The purpose is not to catch the patient out. It is to observe how an account is generated, qualified and revised. Begin with open, neutral questions, then clarify time

and setting. Compare the response with reliable collateral information. Finally, offer the verified context calmly and note whether the patient incorporates it, remains uncertain or resists it.

- Record whether the account was volunteered or followed a direct memory question.
- Describe its content as ordinary and plausible, or fantastic and dream-like, without using those adjectives as diagnoses by themselves.
- Document the response to correction, including reflection, acceptance, persistence or irritation.
- Integrate the observation with the amnestic state and lesion context rather than inferring anatomy from a single reply.

The interview should minimise suggestion. Closed alternatives may insert content that later sounds like the patient's own recollection. Repeated challenge may produce defensiveness that resembles resistance to correction but reflects the interpersonal method rather than the memory disorder. The clean observation is obtained when the examiner states the verified context once, respectfully, and watches whether belief revision occurs.

Collateral verification is indispensable because plausibility can mislead in either direction. A prosaic story may be false, and an unusual story may occasionally be true. Clinical classification must not depend on the examiner's intuition about what "sounds possible." Documentation should separate observed speech from verified facts and from interpretation. This preserves the chain of reasoning for colleagues and prevents a colourful narrative from becoming a diagnostic shortcut.

> **GOOD TO REMEMBER**
>
> Use correction as a bedside probe, not as a contest. The clinically useful observation is whether contextual evidence modifies the account. Escalating confrontation obscures that signal and can turn an assessment of memory monitoring into an assessment of the interview relationship.

19.8 Integration, differential framing and answer construction

Confabulation should be placed within the amnestic formulation before broader psychiatric explanations are invoked. Delirium, primary psychotic phenomena, the dementia syndrome, deliberate fabrication and ordinary memory error may all enter the differential, but their full diagnostic criteria belong to their respective sections and chapters. The dementia syndrome and its workup are covered in the Dementias and neurocognitive disorders notes. Delirium criteria and subtypes are covered in the dedicated delirium section of this chapter.

The candidate's answer can proceed in a compact causal sequence. Define confabulation descriptively as an inaccurate account offered in a memory context. Contrast provoked and spontaneous forms using prompting, content and correction. State the lesion associations. Then explain temporal-context and source-memory failure, adding ventromedial frontal monitoring failure for the spontaneous form. Finish by showing how collateral verification and response to correction convert the theory into bedside assessment.

TRAP **Do not equate confabulation with a confidently told falsehood.** Confidence, plausibility and factual error do not classify the phenomenon. The provoked-spontaneous distinction requires the production pattern, correction response and anatomical context.

The final conceptual distinction is between production and checking. Memory-network damage makes accurate reconstruction difficult; frontal-contextual dysfunction makes an erroneous reconstruction harder to reject. Provoked confabulation is usually exposed by a question and remains revisable. Spontaneous confabulation emerges more independently, becomes fantastic or dream-like, and resists correction. This is why the second pattern has a narrower basal forebrain and ventromedial frontal association.

Provoked confabulation is prompted, prosaic and readily revised; spontaneous confabulation is often unprompted, fantastic, correction-resistant and linked to bilateral basal forebrain with ventromedial frontal dysfunction.

20 · Other causes of amnestic syndrome and the transient amnesias

The useful distinction is anatomical first, temporal second. A persistent, disproportionate failure of episodic learning points toward damage within the medial temporal, diencephalic, basal forebrain or retrosplenial memory network; a self-limited attack demands separation of transient global, epileptic and ischaemic amnesia. The examination marks lie in connecting the memory phenotype to lesion locus and course, not in writing an undifferentiated list of causes.

Chronic amnesic syndrome describes enduring impairment in acquiring or retrieving episodic memories that is prominent enough to require anatomical explanation. It is not synonymous with delirium, because stable amnesia can remain after any acute confusional state has cleared. It is also not synonymous with the dementia syndrome, which is broader and is covered with its full workup in the Dementias and neurocognitive disorders notes. This section instead asks a narrower question: which memory-system lesion best explains the pattern, and when does an apparently isolated attack require a seizure or vascular formulation?

5.2 TO 10 CASES PER 100,000 PER YEAR	NOT OVER 24 HOURS	MINUTES; UNDER 1 HOUR
TGA INCIDENCE	TGA DEFINITION	TEA USUALLY

20.1 A lesion-locus map of persistent amnesia

The network approach prevents the familiar error of treating memory as a function of the hippocampus alone. Persistent episodic-memory impairment may follow hippocampal, diencephalic or retrosplenial damage, while bilateral injury to the medial temporal lobe, diencephalon or basal forebrain can produce a profound amnesic syndrome. These regions are linked components of an extended limbic memory system. A lesion at a node or connecting pathway can therefore create a similar headline complaint while leaving a different retrograde profile and different accompanying signs.

PRINCIPAL LOCUS	GROUNDED CAUSES	LOCALISING EMPHASIS
Hippocampal or medial temporal	Herpes simplex encephalitis, limbic encephalitis, anoxia and temporal-lobe removal	Failure of durable new episodic learning
Diencephalic or basal forebrain	Bilateral thalamic infarction, lesions around the third ventricle and post-subarachnoid-haemorrhage injury	Encoding weakness may coexist with defective retrieval
Retrosplenial	Tumour or haemorrhage	Posterior network disruption can also impair episodic memory

This is a localisation framework, not a claim that all causes are clinically interchangeable. Establish whether registration and attention are adequate, demonstrate failure across a delay, examine the extent of retrograde loss, and use the neurological course and imaging to decide which network component is implicated. The full ordered bedside cognitive examination and standardised instruments belong in the Psychotherapies, clinical assessment, MSE and nosology notes.

Encoding and retrieval language adds precision to that map. A weak delayed response may reflect a memory trace that was never laid down effectively, difficulty gaining access to a stored trace, or both. Diencephalic amnesia can combine weak encoding with defective retrieval; the hippocampal pattern is more closely centred on failure to consolidate new episodic material. This distinction should guide the examiner's explanation, but it should not be overclaimed from a single bedside task. Converging evidence comes from adequate initial attention, performance after a delay, the benefit or absence of benefit from retrieval support, the temporal extent of retrograde loss, collateral history and lesion-compatible imaging. The result is a neurobehavioural formulation rather than a score-labelled complaint.

20.2 Medial temporal injury: infection, autoimmunity and hypoxia

Hippocampal amnesia is classically dominated by impaired anterograde learning with a comparatively limited retrograde deficit. Autobiographical recall may be affected mainly for the most recent **2 to 5 years**, rather than across much of adult life. The practical lesson is that a patient may lose the ability to establish new episodic memories while retaining a large body of older personal knowledge. Immediate conversational coherence does not refute this formulation because immediate maintenance and durable consolidation are different operations.

Herpes simplex encephalitis preferentially injures the medial temporal lobes, including the hippocampus, and produces haemorrhagic necrosis; magnetic resonance imaging may show bilateral temporal-pole T2 hyperintensity. The well-known case of Clive Wearing illustrates the resulting dense anterograde amnesia. The point of the vignette is anatomical: destruction of medial temporal memory structures can leave each present moment poorly connected to the next, even when other aspects of interaction remain recognisable.

Autoimmune disease reaches the same functional system by another route. In anti-NMDA-receptor encephalitis, autoantibodies target the GluN1 subunit of the receptor, a subunit essential

for hippocampal memory consolidation. Cognitive dysfunction therefore belongs within a wider neuropsychiatric presentation rather than being dismissed as poor engagement. Hypoxic-ischaemic injury, whether arising from respiratory, cardiovascular or metabolic compromise or carbon monoxide intoxication, selectively threatens the CA1 hippocampal region. Attempted hanging, cardiorespiratory arrest, carbon monoxide intoxication and inhalational anaesthesia are documented settings in which chronic amnesia may dominate the post-hypoxic picture. The shared endpoint is hippocampal vulnerability, but the causal history remains indispensable.

20.3 Diencephalic, basal forebrain and vascular amnesia

Basal forebrain amnesia may follow rupture of an anterior communicating artery aneurysm, the associated subarachnoid haemorrhage or surgical clipping. The mechanism is injury to small perforating arteries that supply basal forebrain structures, including the fornices. The anatomical label "basal forebrain" should be unpacked rather than left vague: the septal nucleus and diagonal band of Broca are named components, and the fornix links this region to the broader episodic-memory network.

Vascular amnesia is similarly site-dependent. Infarction involving the thalamus, medial temporal lobe or fornices may produce amnesia. A thalamic infarct may initially present with delirium, but the enduring memory deficit must be assessed after attention and arousal have stabilised. This protects against describing a fluctuating failure to register information as if it were a fixed storage disorder. Delirium criteria, mechanisms and management are taught in the dedicated delirium sections of this chapter.

Diencephalic amnesia, exemplified by Korsakoff syndrome, has a different retrograde contour from the relatively focal hippocampal pattern. It can extend back across decades while showing relative sparing of more distant memories, creating a temporal gradient; both encoding and retrieval may be defective. This contrast is a discriminator, not a complete account of Korsakoff pathology. Mammillary-body and thalamic pathology, nutritional antecedents and confabulation are covered in their dedicated sections. For the present differential, extensive temporally graded retrograde loss favours diencephalic network dysfunction, whereas a shorter recent retrograde loss fits hippocampal amnesia more closely.

20.4 The transient-amnesia problem

A transient memory gap is a syndrome description, not a diagnosis. The clinician must reconstruct the episode from a witness wherever possible because the patient may be unable to encode the very events being investigated. The high-yield axes are onset, duration, recurrence, state on waking, preservation of personal identity, attention and consciousness, associated neurological or epileptic phenomena, repetitive questioning, and the memory left after recovery.

The major distinction is between **transient global amnesia**, which is prolonged but self-limited and usually isolated, and transient epileptic amnesia, which is brief and recurrent. A vascular transient amnesia is briefer still: a transient ischaemic attack involving the thalamus or medial temporal lobe may produce 15 to 20 minutes of amnesia. Duration alone cannot make the diagnosis, but it sharply changes the prior formulation.

TGA is uncommon. Its reported incidence is 5.2 to 10 cases per 100,000 people per year. It occurs mainly in later life, peaks around the sixth and seventh decades, and is exceptional below 40 years of age. TEA is usually seen in a similar age range but can occur in younger people; no numerical incidence or age-of-onset estimate is stated here. Epidemiology therefore supports context but must not substitute for the observed attack pattern.

RULE **Classify transient amnesia by the whole episode.** Duration, recurrence, waking relationship, identity, attention, questioning, seizure history and the residual memory gap belong in one reconstruction; no isolated feature is a sufficient label.

20.5 Transient global amnesia: the clinical phenotype

The fundamental event is abrupt inability to retain new factual information. **Anterograde amnesia** is more profound than the accompanying retrograde loss, which chiefly affects recently acquired information. General knowledge and fundamental personal information remain available, so patients can ordinarily state who they are and supply familiar biographical facts. This preservation of identity is diagnostically valuable: the patient is bewildered by the memory failure, not emptied of a personal self.

Working memory remains available, but new material disappears after only a very short interval. The patient is disoriented and repeatedly cycles through the same questions, while consciousness, attention, language and visuospatial function remain intact. Repetitive questioning is therefore not mere anxiety. It is the behavioural signature of repeated registration without durable retention. The preserved attention and language distinguish the episode from a global confusional state.

An operational bedside synthesis, without attaching an unsupported named criteria set, is:

- acute, prominent anterograde amnesia with a less marked recent retrograde component;
- preserved personal identity and general knowledge;
- repetitive questioning despite preserved working memory;
- no clouding of consciousness or impairment of attention;
- no accompanying language or visuospatial deficit;
- complete resolution of the acute syndrome within the defining time ceiling.

Frightening events, physical exertion and sexual activity may immediately precede an attack; stressful emotional or physical events and migraine are also associated. These are contextual clues, not diagnostic proofs. The observed cognitive pattern and course remain primary.

20.6 TGA course, prognosis and disputed mechanism

TGA lasts **3 to 24 hours** and is most severe during the initial 1 to 2 hours; by definition, the acute attack must not exceed 24 hours. Another account places typical restoration of new learning at 4 to 8 hours. Recovery is gradual: retrograde loss contracts as anterograde learning returns, leaving a dense gap for the attack and often for the couple of hours immediately before

it. The patient may therefore look recovered while remaining unable to remember the episode itself.

The aetiology remains unsettled. Proposed mechanisms include transient posterior-cerebral-artery ischaemia affecting medial temporal structures, migraine, seizure, and impaired cerebral venous outflow. Different mechanisms may operate in different patients. Yet the vascular hypothesis must not be promoted to established fact: another account argues that thromboembolic cerebrovascular disease does not explain the usual syndrome, and TGA itself is not a risk factor for stroke. Acute focal or atypical presentations still require a vascular differential; that clinical precaution is separate from claiming that classic TGA is a transient ischaemic attack.

TRAP **Do not flatten the source discrepancies.** The exam-tested acute definition is no more than 24 hours, supported by the account that typical recovery occurs within 4 to 8 hours; a major psychiatry text instead uses the outlying phrase “improving over the course of several days.” Recurrence is also framed differently: one source gives approximately **10% overall**, while another gives **2 to 3% per year**. Preserve “overall” versus “per year.” Both formulations convey low recurrence and need not be forced into one figure.

The general prognosis is excellent. Reassurance is appropriate after the attack pattern has been established and competing urgent causes have been considered. Recurrent, very brief or waking-related events should reopen the diagnosis rather than being repeatedly called TGA.

20.7 Transient epileptic amnesia

Transient epileptic amnesia resembles TGA because confusion, disorientation and repetitive questioning can occur, but its temporal signature is different. Attacks are recurrent, generally last less than 1 hour and commonly only a few minutes; they often occur immediately on waking or after a catnap. Partial recollection of the ictus may survive, and the episodes may evolve into complex partial seizures. Seizure classification and anticonvulsant pharmacology belong in the Epilepsy and psychiatry notes.

The disorder is not confined to attacks. Patients may discover large autobiographical gaps when they try to retrieve holidays or family events from preceding years. A further hallmark is **accelerated long-term forgetting**: new material can be recalled normally after a short delay measured in hours or days, yet be lost after a delay of weeks. Standard short-delay testing can therefore appear reassuring while the patient describes a genuine long-interval failure. This discrepancy should prompt testing and collateral history across ecologically meaningful intervals rather than an inference of poor effort.

A separate source uses **pure epileptic amnesia** for a partial seizure consisting only of anterograde and variable retrograde amnesia. It emphasises paroxysmal onset, brevity, absence of anxious questioning and a history of complex partial or grand mal seizures. The terminology and questioning description are not perfectly identical across sources, so the safer synthesis is clinical: brief recurrent attacks with seizure evidence favour an epileptic mechanism, whether or not repetitive questions occurred.

MNEMONIC

GLOBAL: LONGER, CLEAR, IDENTITY KEPT. EPILEPTIC: BRIEF, RECURRENT, WAKING.

In TGA, the attack lasts longer while attention and identity remain clear. In TEA, brevity, recurrence and onset after sleep carry more discriminating weight.

PITFALL

A normal routine EEG does not exclude TEA. Routine recordings may be normal, whereas sleep EEG usually reveals temporal-lobe spikes or sharp waves. If the history is compelling, the state sampled by the recording matters.

20.8 Assessment, investigation and management

Locus. A patient with active unexplained transient amnesia belongs in an emergency or acute medical setting until dangerous neurological and metabolic alternatives have been addressed. Once stable, outpatient review is appropriate for reconstructing recurrence, long-interval autobiographical loss and sleep-related attacks. Persistent amnesia requires the care setting dictated by its underlying encephalitic, hypoxic, vascular or structural cause.

Focus. During the attack, establish attention, consciousness, identity, language, visuospatial function, focal neurological findings and the ability to retain new material. In the short term, obtain a precise witness account and define onset, offset, precipitant, recurrence, waking relationship, questioning and residual gap. Over longer follow-up, ask specifically about autobiographical holes and accelerated forgetting. Thorough history, neuropsychological evaluation and EEG recordings are required to distinguish TGA from TEA.

Construct the history as a timeline rather than asking whether the patient “was confused.” Ask what the patient was doing immediately before the gap, which question was repeated, whether answers were understood, whether identity and familiar knowledge were retained, whether ordinary speech and navigation were preserved, and what the first durable memory after the event was. Then ask separately about earlier similar spells, sleep or waking, exertion or emotional stress, seizure phenomena and vascular symptoms. This sequence reduces retrospective blending of the attack with the distress around it. It also prevents a normal examination after recovery from erasing an abnormal witnessed event. For persistent amnesia, interview collateral informants about day-to-day learning and remote autobiographical loss, because apparent social fluency may conceal repeated forgetting.

PATTERN	DURATION AND RECURRENCE	POSITIVE CLUES	AFTER THE EVENT
TGA	3 to 24 hours; recurrence is low	Identity and general knowledge retained; attention clear, repetitive questions	Dense gap for the attack
TEA	Under 1 hour, often minutes; recurrent	Often on waking, sometimes partial ictal recall; sleep EEG abnormalities	Autobiographical gaps and accelerated forgetting may persist
Thalamic or medial-temporal TIA	15 to 20 minutes	Vascular localisation must be sought	Do not relabel a vascular event as classic TGA

Modus. Treatment follows the mechanism. Persistent amnesia requires treatment of the underlying medical or structural disease and rehabilitation directed at functional memory problems. For classic TGA, the expected course justifies explanation and follow-up after urgent mimics have been excluded; do not infer a disease-specific drug regimen from the diagnostic label. For TEA, anticonvulsant treatment typically prevents further attacks but may not improve established remote-memory loss. Drug selection belongs with current guidance and the Epilepsy and psychiatry notes.

GOOD TO REMEMBER

Ask beyond the attack. Recurrent waking episodes suggest TEA, but remote autobiographical gaps and accelerated long-term forgetting may be the complaints that expose it. Pair collateral history with sleep EEG when routine recording is unrevealing.

TGA is a longer, self-limited amnesia with preserved identity and attention and a strict 24-hour ceiling; TEA is brief, recurrent, often waking-related and may require sleep EEG to reveal temporal epileptiform activity.

The bedside examination and the focal lobe syndromes

21 · Frontal lobe syndromes: dorsolateral, orbitofrontal and medial subtypes

Frontal lobe syndromes are disorders of organised behaviour, not merely failures on a cognitive screen. The examiner is trying to recognise which aspect of goal-directed conduct has broken down: planning and mental flexibility, inhibition and social regulation, or initiation and drive. This framing matters in psychiatry because the presentation may resemble depression, mania, a personality disorder or a global neurocognitive syndrome even when the primary disturbance is focal prefrontal dysfunction.

The traditional clinical framework describes at least **three** prefrontal syndromes: dorsolateral or frontal-convexity, orbitofrontal or inferior-ventral, and medial frontal or anterior cingulate. They overlap considerably and commonly occur in combination rather than as discrete entities. The

useful answer therefore identifies a dominant pattern, states the expected accompanying signs, and then acknowledges mixed presentations. These syndromes are usually described with bilateral or large unilateral prefrontal damage, so a neat behavioural label provides insufficient evidence for a tiny, isolated lesion.

21.1 The clinical map: plan, restrain and initiate

A practical way to enter the topic is to ask what the patient cannot regulate. The dorsolateral system is expressed clinically through organisation, problem solving and flexible control of a mental set. The orbitofrontal system is expressed through inhibition, appraisal of social context and restraint of behaviour. The medial frontal and anterior cingulate system is expressed through initiation, sustained engagement and the conversion of intention into action. These are clinical organising ideas, not claims that any complex human behaviour resides in a single cortical patch.

PLAN	RESTRAIN	INITIATE
DORSOLATERAL	ORBITOFRONTAL	MEDIAL

The accompanying figure should be read as a map of dominant clinical signatures. It helps organise recall, but the patient in front of you may show apathy together with disinhibition, or executive rigidity together with poor initiation. Such mixtures are expected because the traditional syndromes overlap. Localisation becomes persuasive when the behavioural pattern, neurological signs and lesion context converge.

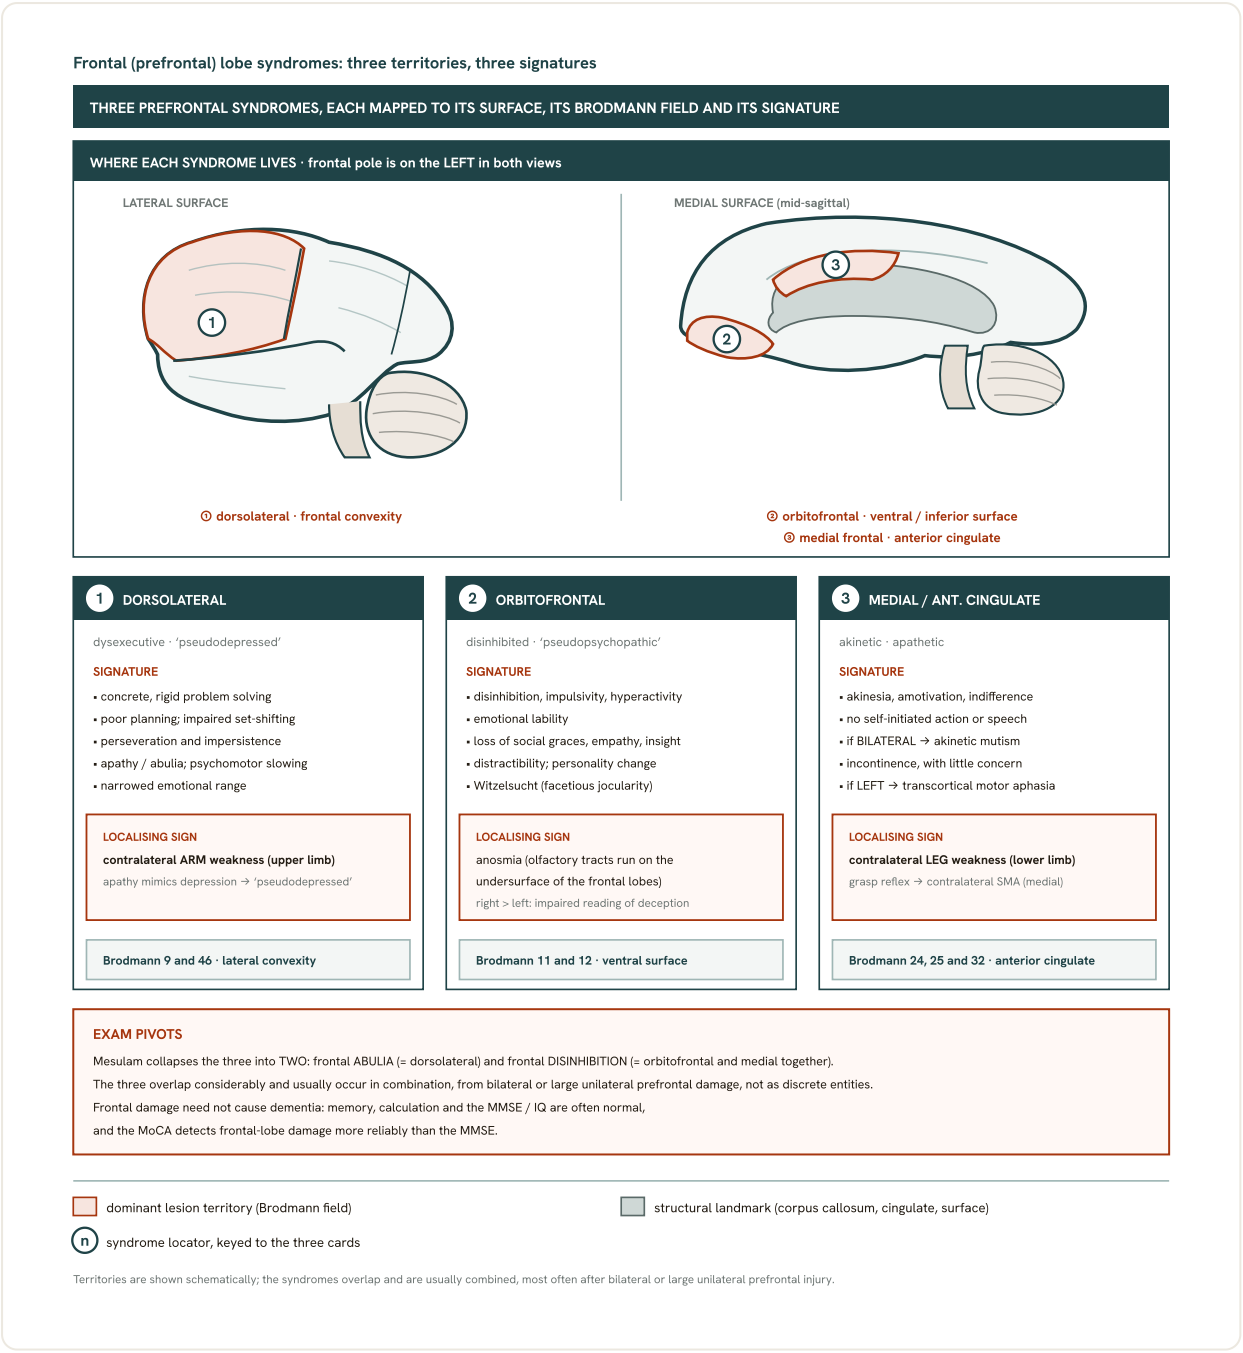


Fig 34.7 **The three frontal-lobe syndromes and their territories.** Dorsolateral (frontal convexity) damage gives a dysexecutive, apathetic picture; orbitofrontal (ventral) damage gives disinhibition and loss of social judgement; medial-frontal and anterior-cingulate damage gives akinesia, apathy and, if bilateral, akinetic mutism. The three overlap and usually occur together (after Kaufman's Clinical Neurology for Psychiatrists and Frontal Lobe Functions).

MNEMONIC

PLAN - RESTRAIN - INITIATE

Plan points to the dorsolateral dysexecutive pattern, **restrain** to orbitofrontal disinhibition, and **initiate** to medial frontal apathy and akinesia. Add the characteristic motor, language, continence and olfactory associations after naming the dominant behavioural change.

RULE **Localise from a convergent syndrome, not from a single behaviour.** Describe the dominant failure, search for the predicted associated signs, and expect overlap when damage is extensive.

21.2 Shared frontal phenomena and behavioural polarity

Across frontal-lobe dysfunction, the patient may become apathetic and indifferent, think and move slowly, struggle to change from one mental set to another, and speak little. **Bradyphrenia** denotes slowed cognition, while **bradykinesia** denotes slowed movement. Paucity of spontaneous speech and action may range from reticence to **abulia**. At the severe end, voluntary movement and speech may disappear together, producing **akinetic mutism**. These features form a low-output pole of frontal dysfunction.

The contrasting pole reflects loss of inhibition. The patient may be distractible, pulled toward whatever is immediately present, socially inappropriate or facetiously jocular. The term **stimulus-bound behaviour** captures environmental capture: conduct is governed by an available cue rather than by an internally maintained goal. This helps explain why a patient can converse fluently yet fail to regulate what is said or done.

These poles are not mutually exclusive. A patient may show sparse self-initiated activity until an external cue triggers an impulsive or inappropriate response. That apparent contradiction becomes intelligible when initiation and inhibition are treated as separable control problems. It is one reason collateral history is essential: ward behaviour, family observations and the patient's functioning outside the interview may reveal abnormalities that a quiet consultation does not.

Observation should distinguish capacity from self-regulation. If the patient completes an action after repeated prompting but does not initiate it independently, the available capacity is greater than the spontaneous output. If the patient understands a boundary yet repeatedly crosses it when a cue appears, comprehension is greater than behavioural restraint. If the patient knows the component steps of an activity but cannot arrange, revise or complete them, stored knowledge is greater than executive organisation. These discrepancies are the clinical texture of a frontal syndrome. They prevent the examiner from reducing all poor performance to memory loss, all inactivity to sadness, or all impropriety to deliberate misconduct.

The tempo and setting of assessment also matter conceptually. A short formal encounter supplies structure, reminders and an explicit goal, all of which can temporarily support a patient whose internally guided behaviour is impaired. Less structured settings may reveal failure to begin, drift from purpose, capture by environmental cues or inability to stop. Collateral examples should therefore be translated into operations: what goal was intended, what cue diverted behaviour, whether correction changed the response, and whether the patient recognised the consequence. This is more localising than a global description such as "personality changed."

The terms apathy, abulia and akinetic mutism should not be used interchangeably. Apathy describes reduced motivation and concern; abulia emphasises marked reduction in spontaneous action and speech; akinetic mutism denotes the extreme combination of absent voluntary movement and absent speech. The medial syndrome supplies the clearest anatomical setting for the severe state, but lower-output features also occur within the dorsolateral pattern.

21.3 Dorsolateral syndrome: the dysexecutive pattern

Damage to the dorsolateral frontal convexity produces the **dysexecutive syndrome**. Problem solving becomes concrete and rigid. Organisational strategies are poor, the patient struggles to shift set, and behaviour may alternate between perseveration and failure to persist. The core problem is not absence of stored knowledge. It is impaired control over how knowledge and action are selected, sequenced, revised and carried through toward a goal.

Perseveration is continued use of a response or strategy after the situation requires change. Impersistence is the complementary failure to maintain an appropriate effort. Together they explain a common bedside impression: the patient may become stuck when flexibility is required, yet abandon a task that needs sustained organisation. Concrete answers and rigid strategies should therefore be interpreted as parts of a control disorder rather than as isolated intellectual deficits.

The affective and motor tempo is often reduced. Apathy or abulia, psychomotor slowing and a restricted emotional range can make the patient look depressed. For this reason, the same phenotype is called the **pseudodepressed syndrome**. The label is a localisation clue, not a claim that depressive illness and frontal dysfunction cannot coexist. Look for the quality of reduced output, executive rigidity and neurological accompaniment rather than deciding from sadness or low activity alone.

Anatomically, the syndrome is associated with Brodmann areas 9 and 46 on the lateral frontal convexity. A lesion may also produce contralateral upper-extremity weakness. The side of weakness matters: it points across the body to the affected hemisphere. When executive disorganisation, psychomotor slowing and this motor pattern occur together, the localisation is stronger than any behavioural feature considered alone.

> **PITFALL**
>
> Do not diagnose primary depression solely from reduced speech, slowed activity and a flat emotional range. Ask whether the patient also shows concrete problem solving, poor organisation, set-shifting failure, perseveration or impersistence, and examine for a contralateral upper-extremity deficit.

21.4 Orbitofrontal syndrome: disinhibition and social dysregulation

Orbitofrontal damage produces the **disinhibited syndrome**, also described as pseudopsychopathic. The patient may become impulsive, hyperactive, distractible and emotionally labile. Social graces, insight and empathy are reduced, so a striking personality change may be more obvious to relatives and staff than to the patient. The deficit is not simply "bad behaviour." It is failure to use context, likely consequences and interpersonal meaning to restrain an immediately available response.

A characteristic manner is **Witzelsucht**: hollow, superficial jocularity with inappropriate joking or laughter. The humour may be socially ill-timed and lacking in genuine emotional depth. An

inappropriate laugh must be interpreted within the wider behavioural context rather than assigned a mechanism from its surface appearance alone.

The lesion lies on the inferior or ventral frontal surface, associated with Brodmann areas 11 and 12, the orbitofrontal cortex proper, with the exact numbering varying by atlas. Poor smell discrimination or **anosmia** may accompany the syndrome because the olfactory nerves lie on the undersurface of the frontal lobes. Smell disturbance is therefore not a decorative extra in the answer. It is an anatomical bridge between the behavioural syndrome and the ventral location.

Laterality can refine the interpretation. Orbitofrontal damage on the right more than the left particularly impairs appreciation of deception.

This association should remain supporting evidence; stand-alone localisation needs the wider syndrome. In practice, extensive lesions commonly blur boundaries, and impaired insight can make self-report reassuring when behaviour has changed profoundly.

GOOD TO REMEMBER

When relatives report a new loss of tact, empathy or restraint, obtain concrete examples and ask about smell. A socially polished interview does not exclude orbitofrontal dysfunction, especially when the patient lacks insight into the change.

21.5 Medial frontal syndrome: failure of initiation

Damage to the medial frontal and anterior cingulate region produces the **akinetic-apathetic syndrome**. The patient is amotivated, apathetic and indifferent, with little self-initiated behaviour or speech. The central clinical question is whether action begins without external prompting. A patient may possess the motor and language capacity required for a response yet fail to generate it spontaneously.

Severity and laterality determine the associated pattern. Bilateral medial frontal damage can produce akinetic mutism, in which the patient is akinetic and mute. A unilateral lesion may cause contralateral lower-extremity weakness, while bilateral damage may weaken both legs. The distribution of weakness is therefore different from the upper-extremity association of the dorsolateral convexity syndrome.

Urinary and bowel incontinence may occur, accompanied by strikingly little concern. The combination of reduced initiation, leg weakness and indifferent incontinence should direct attention toward the medial surface rather than being dismissed as poor cooperation or a nonspecific behavioural problem. Each element reinforces the same localisation when it appears within a coherent syndrome.

Language provides an additional lateralising clue. A left medial frontal lesion involving the anterior cingulate can produce transcortical motor aphasia. The syndrome is anatomically associated with Brodmann areas 24, 25 and 32 in the anterior cingulate gyrus. Extension into the corpus callosum may add alien hand; the disconnection syndrome itself is taught in the dedicated section on alien hand, disconnection and visual association syndromes.

The behavioural description must remain distinct from weakness, aphasia or reduced consciousness. A patient who does not act may be unable, may not understand, may be insufficiently awake, or may lack initiation. The medial frontal formulation becomes credible only after those alternatives are considered and the full pattern aligns.

21.6 Comparison at the bedside

The examination answer is strongest when the syndromes are compared on shared axes rather than presented as separate lists. Start with the dominant control failure, then add behavioural tone, neurological accompaniment and surface anatomy. The table deliberately separates signature from supporting signs so that a single finding is not mistaken for the whole syndrome.

AXIS	DORSOLATERAL	ORBITOFRONTAL	MEDIAL FRONTAL / ANTERIOR CINGULATE
Dominant failure	Planning, organisation, set shifting and persistence	Inhibition, impulse control and social regulation	Initiation of behaviour, movement and speech
Characteristic behaviour	Concrete, rigid problem solving with perseveration or impersistence	Impulsivity, distractibility, emotional lability and social impropriety	Apathy, indifference, amotivation and scant spontaneous output
Memorable descriptive label	Dysexecutive or pseudodepressed	Disinhibited or pseudopsychopathic	Akinetic-apathetic
Motor association	Contralateral upper-extremity weakness	No defining motor association in the permitted account	Contralateral leg weakness, or bilateral leg weakness with bilateral damage
Additional clue	Reduced emotional range and psychomotor slowing	Anosmia and hollow jocularity	Indifferent incontinence; akinetic mutism with bilateral damage
Anatomical surface	Lateral frontal convexity	Inferior or ventral frontal surface	Anterior cingulate on the medial surface

A compact clinical sequence is useful:

- Identify whether the dominant failure concerns planning, restraint or initiation.
- Seek the predicted motor, olfactory, continence, language and affective accompaniments.
- State whether the observed picture is a dominant subtype or a mixed frontal syndrome.

The sequence protects against paired errors. The first is overlocalisation from a colourful isolated behaviour. The second is underrecognition because the patient does not reproduce a complete textbook phenotype. A syndrome is a pattern with internal anatomical coherence, and overlap is part of that pattern rather than evidence against it. Documentation should name the dominant pattern and retain discordant findings, because mixed damage can make those discrepancies clinically meaningful. The final formulation should explain apparent contradictions instead of forcing every finding into the dominant label, and should state which signs remain unexplained.

21.7 Why routine global testing may mislead

A normal global score does not neutralise a compelling frontal history. Frontal-lobe damage does not necessarily cause dementia. Memory, simple calculation and visuospatial perception may remain relatively preserved because they depend substantially on parietal, temporal or distributed systems. Normal results on IQ tests or the MMSE do not exclude frontal damage.

This dissociation explains why the interview must examine behaviour in context. A patient may answer factual questions, calculate simply and copy visuospatial material, yet be unable to organise a plan, shift strategy, inhibit an inappropriate response or initiate purposeful activity. The deficit lies in governing cognition and conduct rather than in the complete loss of the component abilities sampled by a global screen.

The danger is particularly important with indolent disease. A slowly growing lesion in the nondominant frontal lobe may become enormous before producing any sign. The statement is not permission to order imaging indiscriminately. It is a warning that apparently preserved conversation and routine cognition cannot, by themselves, exclude frontal pathology when personality, initiative or executive conduct has changed.

Detailed administration, scoring and interpretation of standardised cognitive instruments belong to the Psychotherapies, clinical assessment, MSE and nosology notes. The bedside cognitive examination, executive testing and frontal release signs are taught in the dedicated section of this chapter. Here the localising lesson is narrower: global preservation and focal behavioural failure can coexist.

TRAP **A normal MMSE does not exclude frontal lobe disease.** Global tests may remain normal because the disabling abnormality is executive, motivational or socially regulatory. History and behavioural observation remain central.

21.8 Alternative model and examination synthesis

The traditional framework is not the only way to organise prefrontal behaviour. The **two**-syndrome model associated with Mesulam describes frontal abulia and frontal disinhibition. Frontal abulia corresponds to the dorsolateral syndrome, while frontal disinhibition combines orbitofrontal and medial-frontal features because both are often present in the same patient.

This alternative should be presented as a teaching compression, not as a contradiction. The more detailed map separates planning, inhibition and initiation because each has a useful dominant anatomical association. The compressed map emphasises the clinically recurring contrast between diminished output and poorly restrained output. Each framework acknowledges that real lesions can produce mixed behaviour.

A high-quality written answer can therefore proceed in a disciplined order. Define frontal syndromes as dominant prefrontal behavioural patterns with substantial overlap. Name the traditional subtypes. For each, state the signature, localisation and associated neurological clue. Contrast the pseudodepressed appearance of dorsolateral damage with the pseudopsychopathic appearance of orbitofrontal damage, then distinguish each from the profound initiation failure of

medial frontal disease. Finish with the alternative model and the warning that global cognition may appear preserved.

At the bedside, avoid moral language. “Lazy,” “rude,” “uncooperative” and “attention-seeking” convert neurological observations into judgements and destroy localisation. Describe initiation, inhibition, set shifting, persistence, emotional range, insight and social regulation in operational terms. Obtain collateral examples and integrate them with the neurological examination. A scan can demonstrate a lesion, but the syndrome is established by the coherent clinical pattern.

Dorsolateral disease impairs planning and set shifting, orbitofrontal disease impairs inhibition and social restraint, and medial frontal disease impairs initiation; mixed syndromes are common.

22 · The bedside cognitive examination, executive testing and frontal release signs

The bedside cognitive examination is a method of clinical localisation, not a ritual list of questions. A good examination asks whether the patient is sufficiently awake to engage, identifies the cognitive process that fails, checks whether a more elementary deficit explains that failure, and then looks for a coherent pattern across interview, testing and behaviour. This matters in psychiatry because fluent conversation, preserved social polish or an acceptable global screen can conceal a disabling problem with initiation, inhibition, set maintenance or self-monitoring.

IN INDIA

Use a brief cognitive test in a version matched to the patient's language and schooling, rather than defaulting to an English instrument.

The marks lie in method. State the prerequisite of arousal, proceed through the domains in an explicit order, choose brief probes that answer a clinical question, and interpret errors rather than merely recording them. The examination should end with executive tasks, behavioural observation, social cognition, judgement, insight and relevant neurological signs. Standardised screens support this process; they do not replace it.

22.1 Establish validity before testing

The bedside cognitive examination begins before a formal question is asked. Observe whether the patient can be engaged, sustain wakefulness, hear and understand the examiner, tolerate the encounter and make an effort that is interpretable. Arousal or consciousness is the prerequisite for valid testing. When that foundation is unstable, an apparent failure of memory, language or executive control may simply reflect an invalid testing state. The proper response is to document the limitation and repeat or adapt the examination, not to construct an elaborate localisation from unreliable performance.

Interpretability also depends on how the task was delivered. Use language the patient can understand, confirm that instructions were grasped, allow for educational and cultural familiarity, and attend to sensory, motor, pain and fatigue barriers. These are not excuses for poor

performance; they are competing explanations that must be tested. Record what the patient actually did: lost the rule, stopped initiating, became stimulus-bound, repeated a previous response, misunderstood the words, or could not execute a motor answer. Such description preserves clinical information that a total score discards.

RULE **Test the process, not just the answer.** First establish arousal and task comprehension; then identify the cognitive operation that failed and ask whether a simpler deficit accounts for it.

Testing is therefore iterative. A surprising error should prompt clarification, a simpler instruction or a non-verbal route where appropriate. Concordance among history, spontaneous behaviour and several bedside probes strengthens an inference. Discordance is useful too: it signals that language, attention, motivation, sensory limitation or task novelty may be influencing the result.

22.2 Use an ordered domain sweep

The examination proceeds from basic engagement toward increasingly integrated functions. After arousal, assess orientation, **attention and concentration**, **working memory**, language, learning and memory, perceptual-motor or visuospatial function, **executive function**, social cognition, and judgement and insight. The order protects interpretation: later tasks often depend on capacities examined earlier.

STAGE	CLINICAL QUESTION	INTERPRETIVE PURPOSE
Arousal	Can the patient remain awake and engage?	Determines whether subsequent performance is valid.
Orientation	Can the patient place the self and situation in context?	Provides a broad anchor, but must be interpreted with attention and memory.
Attention and concentration	Can a mental set be selected and sustained?	Checks the gateway needed by most later tasks.
Working memory	Can information be held while it is manipulated?	Links attention with goal-directed performance.
Language	Can the patient understand and formulate meaningful responses?	Prevents language failure being mislabeled as global cognitive failure.
Learning and memory	Can information be acquired, retained and accessed?	Separates acquisition difficulty from later retrieval difficulty by observing the process.
Perceptual-motor and visuospatial function	Can visual information guide organised action?	Identifies whether a visually mediated task is suitable for executive interpretation.
Executive function	Can behaviour be organised, shifted and inhibited?	Tests control over action rather than stored knowledge alone.
Social cognition, judgement and insight	How is cognition used in interpersonal and real-world decisions?	Reconnects test performance with everyday function and self-appraisal.

Standardised batteries may be comprehensive, but their length can consume time and fatigue the patient. A brief, hypothesis-led bedside examination combined with a suitable screening tool is often the practical clinical approach recommended for the resident. Brevity should come from choosing discriminating probes, not from omitting domains.

AROUSAL	DOMAIN	ERROR	PATTERN
VALIDITY	PROCESS	INTERPRETATION	LOCALISATION

22.3 Attention, concentration and working memory

Attention is not a single observable behaviour. At the bedside, ask whether the patient can register the rule, maintain it, resist distraction, detect a change and continue until the task ends. Concentration is inferred from sustained goal-directed performance. Working memory adds an active demand: information must remain available while the patient reorders, updates or acts upon it. These capacities are best judged from accuracy together with the manner of responding.

Useful probes include:

- digit span forward and backward, comparing simple retention with the additional demand of reversing the sequence;
- months backward, which requires maintenance of a familiar ordered sequence while proceeding against its usual direction;
- alphanumeric sequencing, which places an explicit manipulation demand on working memory; and
- observation across the interview for losing the question, drifting from the task, requiring repeated prompts or abandoning an effort prematurely.

Do not interpret a failed complex task before checking these gateway functions. Poor attention can impair encoding, make recall inconsistent, disrupt comprehension and produce disorganised construction. Conversely, a patient may remain socially engaged yet fail when the task requires holding and manipulating information. This is why a conversational impression is insufficient.

The examiner should note whether errors are immediate or emerge with continued demand, whether the rule is forgotten or merely violated, whether cueing restores performance, and whether the patient recognises the lapse. These observations convert a pass-or-fail exercise into a cognitive formulation. They also guide the next probe: simplify the task when comprehension is uncertain, reduce perceptual or motor demand when output is constrained, and repeat under better conditions when fatigue has made the result uninterpretable.

Documentation should preserve this reasoning. Begin with the testing conditions and degree of engagement, then give a domain-by-domain account of performance and the characteristic errors. Separate what was observed from what is inferred. For example, “lost the alternating rule despite restatement” is an observation; “impaired set maintenance” is an interpretation. Note adaptations, prompting and self-correction because they change what a response means. End with a synthesis that identifies the strongest pattern, the important confounders and the findings that do not fit. This style is more clinically portable than a string of scores: another examiner can

understand the basis of the conclusion, repeat the crucial probes and judge whether change reflects cognition, context or testing conditions. It also guards against circular reasoning, in which a presumed diagnosis dictates the meaning of every error. Bedside examination generates a localising hypothesis; history, neurological examination and appropriate investigation must test that hypothesis.

22.4 Language, memory and visually mediated probes

Domain-specific tools are chosen to clarify a question raised by the bedside examination, not to decorate it. For language, the bedside Western Aphasia Battery-Revised and the Boston Naming Test are named clinical options. For memory, the Hopkins Verbal Learning Test provides a verbal learning probe. Their item content and scoring are not reproduced here. The clinician still has to decide whether poor performance reflects language comprehension, inadequate attention, failed acquisition, retrieval difficulty or an unsuitable testing context.

Visually mediated tasks broaden the examination beyond verbal interaction. Line bisection and target cancellation are bedside probes for neglect. The clock-drawing test and Navon figures probe visuospatial function. A poor result is not self-interpreting: visual perception, spatial organisation, motor output, instruction comprehension and executive planning may all affect the finished response. Describe the error pattern and test the relevant components before assigning a lobe syndrome.

This is also the point at which examination boundaries matter. Hemispatial neglect, apraxia and agnosia require the dedicated assessment taught with the parietal lobe syndromes. Temporal and occipital syndromes likewise need their own focal examination. Here, these bedside probes serve the general method: they reveal which domain needs expansion and prevent a verbal screening encounter from being mistaken for a complete cognitive examination.

> **PITFALL**
>
> A malformed drawing is often labeled “executive dysfunction” at sight. First ask whether the patient saw the material, understood the instruction and could perform the motor act. Executive interpretation becomes credible only after the task’s perceptual, language and output requirements are accounted for.

22.5 Standardised screens: know what each samples

A screening instrument is a structured sample, not the whole examination. Its value depends on the domains it includes, the clinical question and the validity of performance. The recognised multi-domain tools overlap, but their emphasis is not identical.

INSTRUMENT	DOMAINS SAMPLED	BEDSIDE IMPLICATION
Mini-Mental State Examination (MMSE)	Orientation, verbal registration and recall, attention, naming and repetition, verbal comprehension, praxis and visuospatial function	Use as a multi-domain screen; do not treat it as a comprehensive executive examination.
Montreal Cognitive Assessment (MoCA)	Orientation, attention, memory, naming, fluency, verbal repetition and visuospatial-executive function	Its domain sampling includes explicit executive content.
Addenbrooke's Cognitive Examination-Revised and current Addenbrooke's Cognitive Examination-III (ACE-R and ACE-III)	Attention and orientation, memory, fluency, language and visuospatial function	Provides a broad domain profile without replacing hypothesis-led bedside testing.

When frontal-lobe damage is the question, the MoCA is more reliable for detection than the MMSE. That relative advantage does not make either instrument a stand-alone localisation test. A discrepant history, behavioural observation or executive probe still requires explanation.

The Frontal Assessment Battery and antisaccades are named tools for frontal executive assessment, but their item content and scoring are intentionally outside this fragment. The standardised instruments themselves, including administration, scoring, interpretation and licensing boundaries, belong to the Psychotherapies, clinical assessment, MSE and nosology notes.

TRAP **Do not equate a global screening result with an executive examination.** State what the selected screen samples, then add bedside probes of sequencing, rule conflict, inhibition, fluency and behaviour according to the clinical question.

22.6 Executive testing at the bedside

Executive function is the control architecture of goal-directed behaviour. It is inferred from how the patient establishes a set, organises a sequence, maintains a rule, shifts when demands change, suppresses a prepotent response, monitors errors and persists to completion. No **single** probe screens every executive dysfunction. The defensible bedside approach is therefore a small set of tasks with different demands, interpreted alongside spontaneous behaviour.

Motor sequencing or Luria tasks examine whether an ordered motor programme can be learned and maintained. Watch for **perseveration**: the patient may continue a previous element or remain trapped in an established response despite correction. Do not reduce the finding to clumsiness. Check that the component movements are physically possible, demonstrate the sequence clearly, and distinguish failure to acquire the programme from failure to sustain it.

The conflicting-instructions tapping task asks the patient to tap once when the examiner taps twice and twice when the examiner taps once. Success requires holding the rule and overriding direct imitation. The go/no-go tapping task asks the patient to tap once when the examiner taps once but not to tap when the examiner taps twice. Here the critical demand is response

inhibition. Record anticipatory tapping, imitation, rule loss, correction after feedback and awareness of error. Similar-looking failures may arise from different points in the control process.

Verbal fluency samples strategic search and self-monitoring. Semantic generation can use an "animals" cue, while phonemic generation can use "words with s". The clinically useful information is not merely output volume. Note whether production starts readily, clusters around a narrow theme, becomes repetitive, breaks the rule or stops despite encouragement. Language capacity and educational familiarity must be considered before calling the pattern executive.

Finally, observe the patient between explicit tasks. **Utilisation behaviour** is suggested when visible objects pull action toward their use without an appropriate goal. Echopraxia and echolalia reveal excessive capture by another person's movement or speech. **Motor impersistence** is difficulty sustaining a requested motor act. These signs may be more revealing than a polished verbal response because they test control in the presence of competing environmental cues.

MNEMONIC

Sequence, Conflict, Stop, Search, Sustain

The executive bedside spine: organise a motor **sequence**, reverse the examiner under **conflict**, **stop** a primed response, strategically **search** for words, and **sustain** a motor set while watching for environmental capture.

22.7 From error pattern to frontal localisation

Localisation should follow the cognitive analysis, not precede it. A patient who fails sequencing because the instruction was not understood has not demonstrated a frontal deficit. A patient who repeatedly understands the rule, can execute its components, yet cannot maintain, shift or inhibit the response provides a more coherent executive pattern. The same principle applies to fluency and working memory: interpret the mechanism of failure and seek convergence across tasks.

The anatomical associations also differ among probes. Go/no-go performance depends on the orbitofrontal cortex. The other executive tasks, including sequencing and fluency, depend chiefly on the dorsolateral prefrontal cortex and its circuit. Motor impersistence is associated with right-hemisphere dysfunction. These are associations for clinical reasoning, not a licence to localise from an isolated error. The accompanying figure places focal bedside signs in a broader lobe and hemisphere framework.

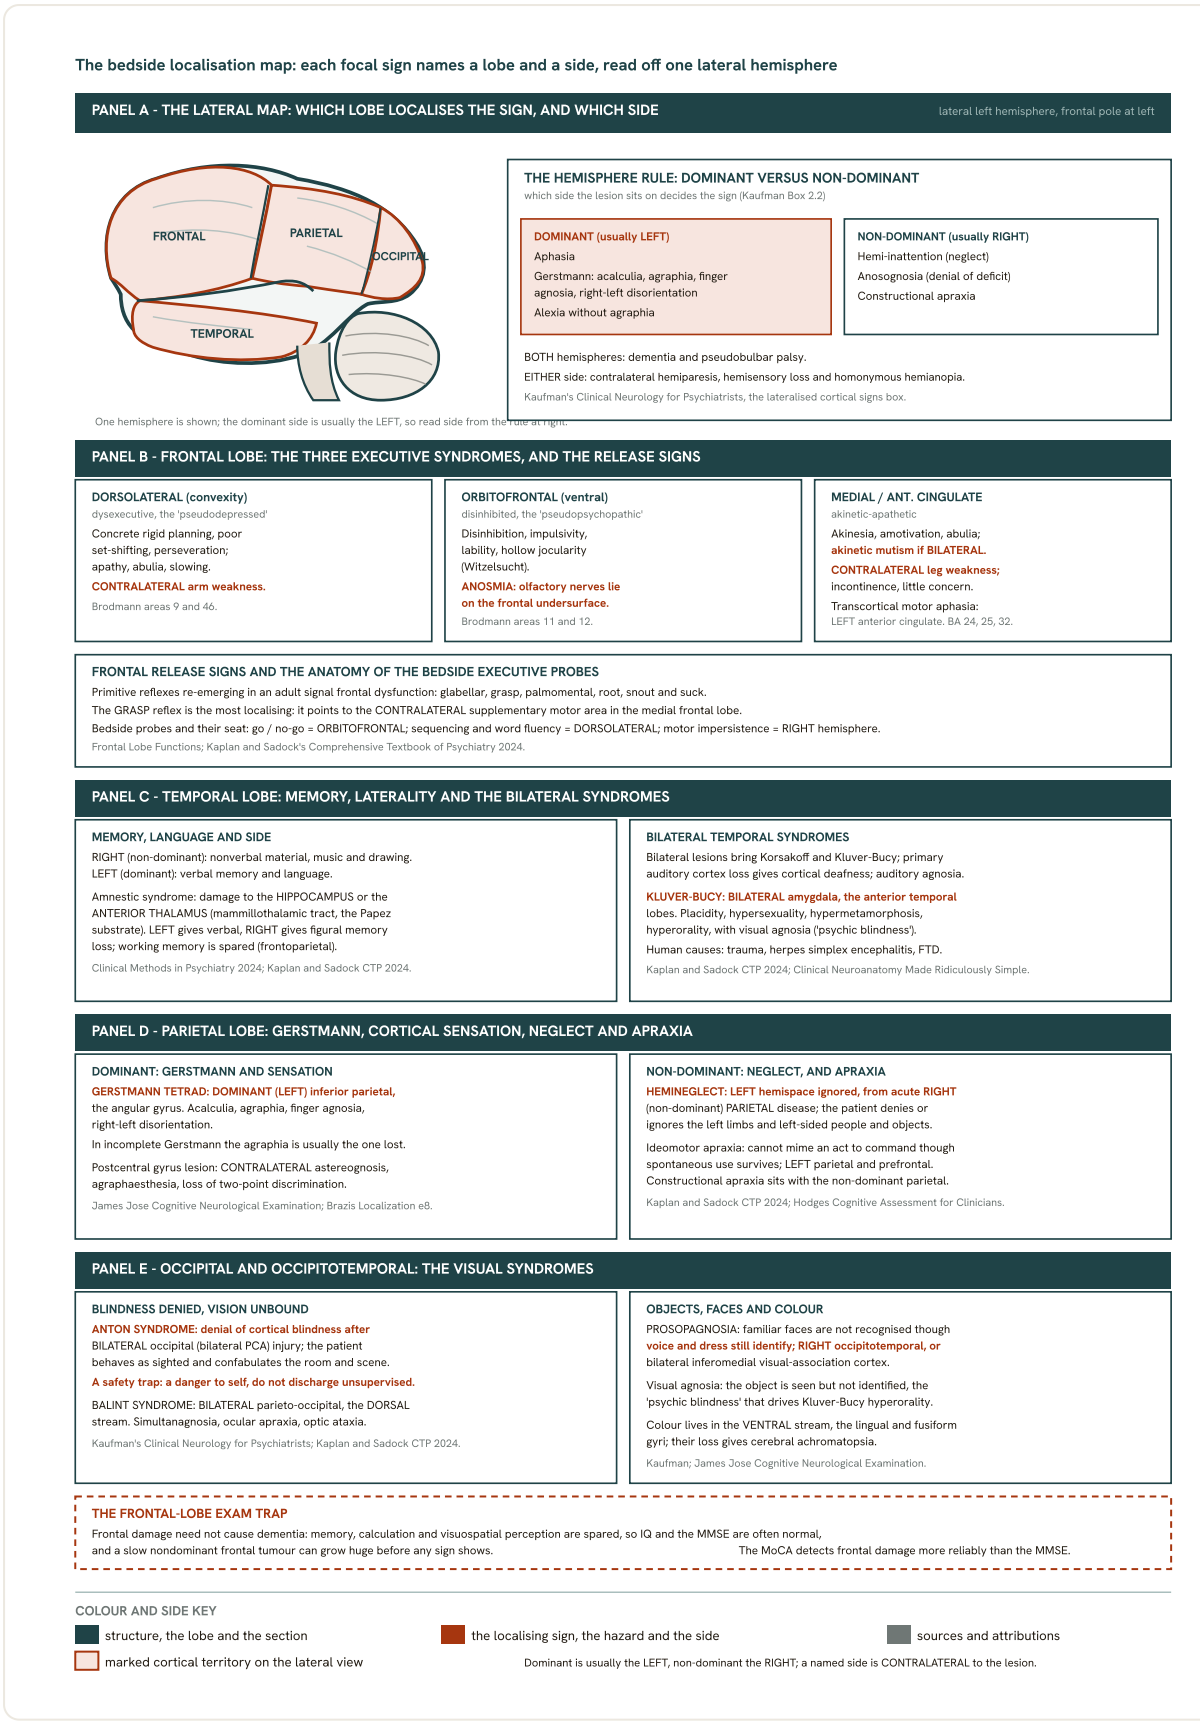


Fig 34.8 **The bedside localisation map.** Each focal sign is read back to a lobe and a side on one lateral hemisphere: frontal for the executive syndromes and release signs, temporal for memory and the Kluver-Bucy syndrome, parietal for the Gerstmann tetrad and neglect, occipital for the visual syndromes; the hemisphere rule fixes dominant against non-dominant. (Kaufman's Clinical Neurology for Psychiatrists; Frontal Lobe Functions; Kaplan and Sadock's Comprehensive Textbook of Psychiatry 2024; Clinical Methods in Psychiatry 2024; James Jose Cognitive Neurological Examination.)

Behavioural observation supplies ecological corroboration. Difficulty initiating the task, stimulus-bound responding, failure to stop, repetition after correction and poor awareness of error may reveal different aspects of impaired control. Their interpretation is strengthened when they agree with the interview and structured probes. The full phenomenology of dorsolateral, orbitofrontal and medial frontal syndromes belongs with the frontal lobe syndromes; this section owns the examination used to detect and localise their cognitive signs.

GOOD TO REMEMBER

If an executive task is abnormal, repeat the instruction in simpler language and verify the elementary motor or language capacity needed for the response. A reproducible failure of rule maintenance or inhibition is more informative than an unexplained wrong answer.

22.8 Frontal release signs

Frontal release signs, also called primitive reflexes, are responses present in infancy and normally suppressed through development. Their re-emergence in an adult suggests frontal-lobe dysfunction. They are most often encountered with acute or severe injury and may be absent when frontal disease is less severe or sub-acute. Presence therefore supports a frontal hypothesis, while absence does not dispose of it.

The bedside repertoire includes the following:

The glabellar response is failure of blinking to extinguish during repeated tapping above the nose. The palmomental response is contraction of the ipsilateral chin or mentalis after stimulation over the thenar region. The snout response is puckering of the lips after the upper lip is tapped. The root, suck and grasp reflexes complete the named group. Elicit them deliberately, explain the manoeuvre, avoid startling the patient, and document the observed movement rather than writing “release signs positive” without qualification.

The **grasp reflex** is particularly useful because it is fairly specific for frontal-lobe injury and carries localising value. Stroke the palm distally toward the base of the fingers and observe for grasping. A grasp response points to dysfunction of the contralateral supplementary motor area in the medial frontal lobe. Laterality is essential: the pathological response is on the side opposite the implicated supplementary motor area.

No **single** frontal release sign is pathognomonic. A response must be interpreted with age and clinical context, the quality and reproducibility of the manoeuvre, associated executive findings and the rest of the neurological examination. A cluster of convergent abnormalities is more persuasive than an isolated equivocal movement. Equally, a normal release-sign examination cannot exclude sub-acute or less-severe frontal disease.

Arousal validates the examination; ordered domains identify the failed process; convergent executive probes and release signs support localisation, but no isolated task or primitive reflex makes the diagnosis.

23 · Temporal lobe syndromes and the Kluver-Bucy syndrome

Temporal-lobe syndromes become manageable when the answer is organised by domain, side and lesion extent. The temporal lobes participate in auditory perception, language, memory and emotion, so their disorders can present to psychiatry as forgetting, altered recognition, disturbed behaviour or an apparent language problem rather than as an obvious motor deficit. The marks lie in showing that these manifestations are related but not interchangeable. A strong formulation names the impaired operation, lateralises it where possible, and asks whether the lesion must be unilateral or bilateral.

> **IN INDIA**
>
> Start aciclovir early on clinical suspicion of herpes simplex encephalitis rather than waiting for laboratory confirmation, because delay worsens the amnestic outcome.

The central contrast is between dominant, non-dominant and bilateral involvement. Dominant temporal dysfunction preferentially disturbs verbal material and language; non-dominant dysfunction preferentially disturbs figural, visual and other nonverbal material; bilateral injury can disrupt auditory recognition or produce a much broader memory-emotion-behaviour syndrome. The most distinctive bilateral behavioural pattern is the **Kluver-Bucy syndrome**, whose feature names are often remembered better than their meanings. The examiner should therefore translate each term into observable behaviour and state the bilateral anterior-temporal or amygdala localisation.

23.1 Functional map and the logic of laterality

The **temporal-lobe syndrome** is not a single uniform phenotype. It is the clinical expression of failure within a lobe that supports auditory perception, language, memory and emotion. The patient may therefore complain of hearing without understanding, remembering some kinds of material better than others, difficulty finding words, or a marked change in emotional and motivational behaviour. These complaints must be converted from everyday language into the specific operation that has failed.

Dominant temporal-lobe dysfunction, usually represented clinically by the left side, interferes with language and verbal memory. The critical word is verbal: difficulty retaining a story, names or spoken material is more lateralising than the global statement that "memory is poor." **Non-dominant temporal-lobe dysfunction**, usually represented by the right side, impairs nonverbal information such as drawing and music. This is not merely a weaker version of the dominant syndrome. The material being processed is different.

LEFT	RIGHT	BILATERAL
VERBAL AND LANGUAGE	FIGURAL AND NONVERBAL	AUDITORY AND LIMBIC SYNDROMES

Laterality is useful only when the task matches the material. A verbal task cannot adequately test the non-dominant temporal contribution, and a figural task cannot by itself establish dominant verbal-memory failure. At the bedside, compare learning across material types, relate the result to

language competence, and seek convergence from history and neurological findings. The aim is not to infer an exact lesion from a single failed item. It is to establish a coherent side-sensitive pattern.

RULE **Localise by operation, material and side.** "Poor memory" is too broad. Verbal learning and word retrieval point toward dominant temporal dysfunction; figural, visual and other nonverbal learning point toward non-dominant dysfunction; broader auditory or behavioural syndromes raise the question of bilateral involvement.

23.2 How to examine a suspected temporal syndrome

A temporal-lobe examination begins before formal testing. Clarify what the patient means by "I cannot remember," "I cannot hear properly," or "words do not come." Determine whether information was understood at presentation, whether it can be maintained briefly, whether it remains available after distraction, and whether the difficulty changes with verbal versus visual material. Ask relatives for concrete examples, because fluent conversation can coexist with failure to retain new events and the patient may not appreciate the pattern accurately.

The examination should sample the owned temporal domains without reproducing a general cognitive battery. Test acquisition and later recall with verbal material, then use nonverbal or figural material to look for a side-sensitive dissociation. Assess ordinary language and word retrieval so that an apparent verbal-memory deficit is not interpreted without considering language performance. Present meaningful environmental sounds and ask what they signify, while separately establishing whether sound is perceived at all. When bilateral anterior-temporal disease is plausible, observe emotional responsiveness, stimulus-driven attention, oral exploration, sexual disinhibition and placidity.

This sequence separates closely neighbouring questions. Perception asks whether the stimulus reaches awareness. Recognition asks whether the percept has meaning. Memory asks whether newly encountered material becomes durably available. Language asks whether verbal material can be understood and expressed. Behavioural observation asks whether emotional and motivational regulation has changed. A patient can fail one operation while retaining another, so a global label such as "organic behaviour" or "cognitive decline" throws away localising information.

Formal test administration and scoring belong to the Psychotherapies, clinical assessment, MSE and nosology notes. The temporal examination uses the same disciplined approach but asks a narrower anatomical question. Repeat an apparently abnormal task in a different form, document whether the material was verbal or nonverbal, and record the type of error rather than only success or failure. A side-sensitive dissociation is more persuasive when language, attention and immediate maintenance have been considered.

GOOD TO REMEMBER

When a patient reports “hearing loss,” ask separately whether sound is detected and whether it is recognised. When the complaint is “memory loss,” ask what kind of material is lost and whether it disappears immediately or only after a delay. Those distinctions turn a vague psychiatric complaint into a localising examination.

23.3 The amnestic syndrome and medial-temporal memory

The **amnestic syndrome** consists of anterograde amnesia for new material together with variable retrograde amnesia. The core clinical problem is impaired formation of durable new memories, not loss of every mental operation. A patient may follow the interview, hold a question in mind and answer coherently, yet later fail to retain the encounter. This apparent paradox is expected because short-term or working memory is spared through its dependence on frontoparietal mechanisms.

Anatomically, damage to the hippocampus can produce the syndrome. The permitted anatomical differential also includes the anterior thalamus and the mammillothalamic tract. That comparison matters because a severe new-learning disorder is not uniquely temporal. Detailed thalamic and mammillary pathology in Korsakoff syndrome belongs in its dedicated section of this chapter; here the useful point is that medial-temporal and diencephalic injury can converge on an amnestic presentation.

The memory deficit is lateralised by material. Left-sided damage produces predominantly verbal memory impairment, whereas right-sided damage produces predominantly figural or visual memory impairment. This is the same material-sensitive logic that organises the broader temporal syndrome. It prevents the common mistake of treating “memory” as anatomically indivisible. In practice, language failure can contaminate verbal learning, and poor initial engagement can mimic defective retention, so the examiner should establish comprehension and adequate acquisition before interpreting delayed recall.

A bedside formulation should state the operation, the material and the preserved capacity. For example, describe impaired acquisition of durable verbal information with intact immediate maintenance, rather than simply recording forgetfulness. If figural learning is selectively poor while verbal learning is relatively stronger, state that contrast explicitly. The full causes, transient amnesias, confabulation and rehabilitation of amnestic disorders are taught in their dedicated sections; the temporal-lobe question is the hippocampal contribution and its laterality.

PITFALL

Do not dismiss medial-temporal amnesia because the patient can repeat information or converse normally for a short interval. Immediate maintenance may be preserved while durable learning is profoundly impaired. Test retention after interference and obtain collateral examples of repeated forgetting.

23.4 Cortical deafness and auditory agnosia

Auditory syndromes demonstrate why perception and recognition must be separated. **Cortical deafness** follows damage to the primary auditory cortex in the setting of bilateral temporal involvement. The ears are structurally intact, yet the person cannot hear. The syndrome therefore localises failure beyond the peripheral organ. It should not be described as inattention merely because the patient does not respond to spoken material.

Auditory agnosia is different: sound recognition is impaired while language remains intact. Sound reaches awareness, but its significance is not identified. The patient's difficulty is therefore not equivalent to deafness, and preserved language argues against explaining the entire problem as a language disorder. At the bedside, first establish whether a sound is perceived, then ask the patient to identify or explain it. This order prevents failure of perception from being mislabelled as failure of recognition.

CLINICAL AXIS	DOMINANT TEMPORAL PATTERN	NON-DOMINANT TEMPORAL PATTERN	BILATERAL TEMPORAL PATTERN
Material most affected	Verbal information and language	Drawing, music and other nonverbal information	Broader failure that cannot be reduced to a unilateral material preference
Memory expression	Predominantly verbal impairment with left-sided damage	Predominantly figural or visual impairment with right-sided damage	Bilateral lesions may produce major amnestic syndromes
Epilepsy-associated cognitive profile	Verbal-memory and word-retrieval deficits	Visuospatial and visual-memory deficits	Classify seizures in the dedicated epilepsy chapter
Auditory consequence	No isolated auditory syndrome is assigned here	No isolated auditory syndrome is assigned here	Cortical deafness after primary auditory-cortex damage, or impaired sound recognition in auditory agnosia
Behavioural consequence	Interpret emotion and behaviour within the whole syndrome	Interpret emotion and behaviour within the whole syndrome	Anterior-temporal or amygdala damage may produce the characteristic placid, oral, sexual and stimulus-bound pattern

The table is a localisation framework, not a substitute for examination. Hearing, recognition, language and memory must each be tested at the level appropriate to the complaint. A bilateral label should be reserved for a pattern whose anatomy and breadth demand bilateral temporal involvement, not applied merely because several symptoms coexist.

23.5 Temporal-lobe epilepsy as a lateralising context

Temporal-lobe epilepsy provides a clinically useful demonstration of the same hemispheric specialisation. In dominant temporal-lobe epilepsy, verbal-memory and word-retrieval deficits

are associated features. In non-dominant temporal-lobe epilepsy, visuospatial functioning and visual memory are preferentially affected. The pattern does not create a separate rule for epilepsy; it confirms that the side and the kind of information remain linked.

The psychiatric assessment should not turn this association into a diagnosis from cognitive testing alone. Word-retrieval difficulty or weak visual memory can support lateralisation within a known clinical context, but neither establishes that episodes are epileptic. Reconstruct the event history, determine whether cognitive problems are persistent or episodic, and seek neurological corroboration. Seizure classification, electroencephalographic interpretation and anticonvulsant pharmacology belong to the Epilepsy and psychiatry notes.

There is also an important acquired-lesion overlap. Herpes simplex encephalitis has a predilection for the temporal lobes and the undersurface of the frontal lobes. When its temporal damage is bilateral, the human behavioural syndrome may coexist with amnesia and complex partial seizures. This cluster should be read anatomically: memory structures, limbic behaviour and seizure susceptibility can be affected within the same regional disease. It does not justify importing the entire classification or treatment of epilepsy into the temporal-lobe answer.

TRAP **Do not equate every temporal cognitive deficit with epilepsy.** Dominant verbal-memory and word-retrieval difficulty, or non-dominant visual-memory difficulty, may lateralise dysfunction in temporal-lobe epilepsy; they do not by themselves prove an epileptic mechanism.

23.6 Kluver-Bucy syndrome: anatomy and feature meanings

The defining anatomical principle is **two-sided** injury. The syndrome depends on damage to the amygdala in the anterior temporal lobes on both sides, producing a mixture of emotional, perceptual and motivational change. A unilateral amygdala lesion should not be offered as the classic localisation. The description arose from the captive-monkey syndrome, but in clinical practice each term must be translated into behaviour rather than recited as an unexplained list.

The cleanly grounded behavioural core contains **four** elements:

- **Tameness or placidity:** reduced aggression and striking docility.
- **Hypersexuality:** excessive and indiscriminate sexual behaviour.
- **Hypermetamorphosis**: compulsive or forced attention to environmental stimuli, so behaviour is repeatedly captured by what is present.
- **Hyperorality**: mouthing of nonfood objects, an abnormal oral exploration of the environment.

A separately grounded associated feature is **visual agnosia**, meaning inability to recognise objects visually despite seeing them. Its anatomical wording requires care: the amygdala lesion accounts for placidity, hypersexuality and hypermetamorphosis, while visual agnosia reflects injury to nearby areas. Thus the perceptual feature belongs in the taught clinical picture, but it should not be simplistically attributed to the amygdala alone.

MNEMONIC

PLACID - SEXUAL - STIMULUS-BOUND - ORAL - SEES BUT DOES NOT KNOW

Placid recalls reduced aggression, **sexual** recalls indiscriminate hypersexuality, **stimulus-bound** unpacks hypermetamorphosis, and **oral** recalls mouthing of nonfood objects; **sees but does not know** captures the separately grounded visual agnosia.

23.7 Human presentations and causes

Humans usually show an incomplete phenotype rather than the full experimental picture. The human Kluver-Bucy syndrome usually occurs in partial form. This is clinically important because insisting on every named manifestation will miss plausible cases, while using one colourful behaviour as proof will overdiagnose them. The formulation should identify the observed features, state what is absent or unknown, and demonstrate compatible bilateral anterior-temporal involvement.

The permitted causes are trauma, herpes simplex encephalitis and frontotemporal dementia. These diagnoses supply different clinical contexts, but the localisation question is shared: can the pathology affect the anterior temporal-amygdala system bilaterally? The syndrome should therefore direct anatomical reasoning without replacing the aetiological assessment.

Herpes simplex encephalitis is especially instructive because it preferentially affects the temporal lobes and the undersurface of the frontal lobes. Bilateral temporal damage can then produce the human syndrome together with amnesia and complex partial seizures. Behaviour that appears disinhibited, sexually inappropriate, orally exploratory or unusually placid must therefore be interpreted alongside memory change and the neurological context rather than assigned prematurely to a primary psychiatric disorder.

Conversely, none of the listed causes can be inferred from the syndrome alone. The behavioural pattern localises a network; the history, examination and appropriate investigation establish the pathology. Human partiality also means that absence of a feature is less decisive than the anatomical coherence of those present. Documentation should avoid claiming a complete syndrome when only a subset is observed.

23.8 Bedside synthesis and examination answer

A concise long-answer structure begins with function, proceeds to laterality, then separates focal domain syndromes from the bilateral behavioural syndrome. State that the temporal lobes support auditory perception, language, memory and emotion. Contrast dominant verbal-language deficits with non-dominant figural and nonverbal deficits. Add hippocampal amnesia with preserved working memory, distinguish cortical deafness from auditory agnosia, and then describe the bilateral amygdala basis and clinical meaning of the characteristic behavioural features.

At the bedside, use a disciplined synthesis. Name the failed operation: perception, recognition, language, durable learning, emotional regulation or stimulus control. State the material involved: verbal, figural, visual, musical or environmental sound. Use side-sensitive findings to distinguish

dominant from non-dominant temporal dysfunction. Reserve the bilateral formulation for cortical deafness, major bilateral memory syndromes, or the anterior-temporal amygdala pattern. Describe each observed behavioural feature in plain clinical terms and record whether the human presentation is partial. Finally, separate localisation from aetiology: the pattern identifies the involved system, while the clinical context identifies trauma, encephalitis, degeneration or another pathology.

The main differentials are operational rather than merely diagnostic. Failure to perceive sound is not failure to recognise it. Failure to recognise a visual object is not blindness. Failure to retain new information after a delay is not disproved by intact immediate repetition. Verbal-memory impairment is not interpretable without language assessment. A striking behavioural change is not enough for the bilateral syndrome unless the feature cluster and anatomy cohere.

Keep ownership boundaries equally clear. Korsakoff syndrome pathology and confabulation are taught in their dedicated sections of this chapter. The general dementia syndrome and its workup belong to the Dementias and neurocognitive disorders notes. Seizure classification and drug selection belong to the Epilepsy and psychiatry notes. These cross-references preserve the focal temporal answer: what operation failed, which side is implicated, whether bilateral damage is required, and how the behavioural terms translate at the bedside.

> Temporal localisation follows material and laterality: left dominant lesions impair verbal memory and language, right non-dominant lesions impair figural and nonverbal memory, and the characteristic placid, hypersexual, stimulus-bound and hyperoral syndrome requires bilateral anterior-temporal amygdala damage.

24 · Parietal lobe syndromes: Gerstmann syndrome, neglect, apraxia and agnosia

Parietal syndromes are best understood as failures of integration. Strength, primary sensation, coordination, comprehension and willingness may be adequate, yet the patient cannot combine symbols, spatial attention, learned action or higher sensory information into an effective response. That contrast is why these syndromes are disproportionately useful at the bedside: a brief task can reveal the impaired operation as well as the likely hemisphere.

The marks lie in disciplined distinctions. The dominant inferior parietal region points toward the Gerstmann cluster; the non-dominant parietal region points toward contralateral hemispatial neglect; praxis must be tested only after elementary motor and sensory failure has been excluded; and cortical sensory loss must be distinguished from loss of primary sensation. A good answer therefore moves from definition, through bedside demonstration, to localisation, while stating what the finding is not.

DOMINANT	NON-DOMINANT	COMMAND	TOUCH
GERSTMANN PATTERN	NEGLECT	IDEOMOTOR FAILURE	CORTICAL SENSORY SIGN

24.1 The parietal map: domain before label

The parietal examination samples several related but separable domains. Symbolic operations include writing, calculation, finger representation and right-left orientation. Spatial attention determines whether the full environment and body enter awareness. Praxis converts an understood intention into a learned act. Higher somatosensory processing converts intact elementary sensation into recognition of form, identity or a traced symbol. A lesion may disturb a particular domain more than the others, so "parietal syndrome" is a localisation hypothesis rather than a uniform behavioural phenotype.

Laterality gives the initial organising axis. Dominant-hemisphere cortical signs include aphasia, the Gerstmann pattern and alexia without agraphia, whereas non-dominant signs include neglect, anosognosia and constructional apraxia. Either hemisphere can also produce contralateral weakness, hemisensory loss or a homonymous field defect. Those elementary findings support a side but do not by themselves identify the higher parietal operation that has failed. Extensive or distributed disease may blur these teaching patterns.

CLINICAL AXIS	DOMINANT PARIETAL EMPHASIS	NON-DOMINANT PARIETAL EMPHASIS	EITHER SIDE / SUPPORTING SIGN
Core higher function	Calculation, writing, finger recognition and right-left orientation	Attention to contralateral hemispace	Interpret as a pattern, not from an isolated task
Characteristic syndrome	Gerstmann syndrome	Hemi-inattention and constructional apraxia	Contralateral motor, sensory or visual-field signs
Common conceptual error	Calling every writing failure aphasic	Calling every left-sided omission a primary sensory loss	Localising from an abnormal response without checking prerequisites

RULE **Define the failed operation, verify its prerequisites, and only then localise.** A higher cortical sign becomes convincing when the patient is sufficiently alert, understands the task, cooperates, and has enough elementary motor, sensory and visual capacity to attempt it.

24.2 Gerstmann syndrome: the tetrad

Gerstmann syndrome is a **four**-component syndrome comprising **acalculia**, **agraphia**, **finger agnosia** and right-left disorientation. Its classic localisation is the dominant, usually left, inferior parietal lobule in the region of the angular gyrus. The value of the tetrad is not merely memorising a list. The components converge on learned symbolic and body-schema operations that are anatomically close enough to fail together.

At the bedside, examine each component directly and document the nature of the error. Calculation should sample more than rote recitation and should separate misunderstanding from faulty manipulation of numerical information. Writing should be spontaneous or elicited in a way that allows inspection of letter formation and organisation, rather than judged from the patient's signature alone. Finger recognition requires the patient to identify or indicate fingers reliably.

Right-left orientation should be tested on the patient's body and interpreted cautiously if language comprehension is doubtful.

MNEMONIC

WRITE - CALCULATE - FINGERS - SIDES

The compact sweep is: ask the patient to write, perform a calculation, identify fingers, and distinguish right from left. The mnemonic preserves the whole tetrad without pretending that an isolated error establishes the syndrome.

The localisation statement should follow the findings, not substitute for them: a coherent tetrad indicates dominant inferior parietal dysfunction near the angular gyrus. If aphasia, weakness, sensory loss, visual impairment or poor attention makes a task uninterpretable, record that limitation rather than scoring the response as another cortical deficit. In an examination answer, this caveat shows clinical maturity because the same apparent failure can arise from different broken prerequisites.

24.3 Gerstmann components can dissociate

The classical tetrad is a pattern, not an all-or-none bedside law. In incomplete Gerstmann syndrome, agraphia is usually the absent component. Therefore, the remaining concordant elements should prompt careful retesting and an anatomically reasoned formulation rather than reflex rejection of the localisation. Conversely, an isolated element should retain a broader explanation until the other components and their prerequisites have been examined.

The nature of the writing disorder is a favourite discriminator. The agraphia in this syndrome is considered apraxic rather than aphasic and need not be accompanied by alexia. The patient's problem is therefore not automatically reducible to loss of language formulation. Examine spoken language and reading separately, and describe how the written output fails. This prevents a visible writing problem from being swept into "aphasia" without analysis.

The same discipline applies to the other elements. Poor calculation may reflect failure to understand the instruction, impaired attention or educational familiarity rather than a focal symbolic deficit. Incorrect finger naming may become uninterpretable when naming itself is impaired. Right-left errors may arise because the patient did not understand the relational language. The examiner should simplify instructions, demonstrate the task when demonstration does not destroy the construct, and seek consistency across repeated responses without turning the encounter into a formal psychometric assessment.

PITFALL

Do not diagnose the full syndrome from an isolated classroom-like error. Establish that the patient understood the instruction, could see and move adequately, and was attending. Then test all components and state whether the pattern is complete, incomplete or only suggestive.

Detailed administration and scoring of standardised cognitive instruments belong to the Psychotherapies, clinical assessment, MSE and nosology notes. The present task is focal localisation: identify the operation that failed, show that the failure is reproducible, and integrate it with the rest of the neurological examination.

24.4 Hemispatial neglect: an attentional disorder

Hemispatial neglect, or hemi-inattention, is a focal disorder of attention. It almost always affects left hemispace in acute right, non-dominant hemisphere disease, most characteristically from a parietal lesion. Distinguishable neglect patterns may also follow frontal or cingulate lesions, so the bedside sign is strongly lateralising but must still be integrated with the wider syndrome.

Gross neglect may be visible before formal testing. The patient may fail to attend to people or objects on the left, ignore the left limbs, or deny that those limbs belong to the self. Observe spontaneous orientation during conversation, approach from each side, and note whether grooming, interaction or response consistently excludes a hemispace. Structured paper-and-pencil or copying tasks may make subtle asymmetry easier to see, but interpretation still depends on adequate vision, motor output and task comprehension.

The decisive conceptual distinction is between inattention and elementary loss. A patient with neglect behaves as though information from a side does not gain sufficient access to ongoing behaviour. A primary sensory or visual deficit, by contrast, concerns the incoming signal itself. The disorders can coexist, which is why the examiner compares visual, tactile and auditory behaviour, checks elementary function, and looks for a consistent spatial bias across modalities and activities. Do not convert every omission into neglect before this groundwork.

Neglect must also not be equated with anosognosia. The former is a spatial-attention disorder owned here; the latter concerns unawareness or denial of deficit and is taught with post-stroke psychosis, emotional incontinence and denial of deficit. They may accompany the same non-dominant hemisphere syndrome, but either term retains its distinct meaning.

> **GOOD TO REMEMBER**
>
> When the patient repeatedly misses material on the left, change the mode of presentation and verify vision, sensation, movement, comprehension and arousal. A cross-modal leftward attentional bias carries more localising weight than an untidy drawing or a missed touch.

24.5 Apraxia: definition and classification

Apraxia is inability to perform complex motor acts despite intact motor and sensory systems, coordination, comprehension and cooperation. These exclusions belong inside the definition because weakness, ataxia, sensory loss, misunderstanding and refusal can each mimic a failed learned action. “Could not do it” is therefore an observation; apraxia is the reasoned conclusion reached after the simpler explanations have been tested.

A useful clinical classification describes **four** principal motor apraxias. The table is a map of dominant associations, not permission to force a patient into a rigid compartment.

TYPE	DOMINANT CLINICAL PROBLEM	LOCALISATION IN THE PERMITTED ACCOUNT	BEDSIDE EMPHASIS
Limb-kinetic apraxia	Execution of a skilled limb act	Basal ganglia or supplementary motor area	Inspect the quality of the learned movement after excluding weakness and incoordination
Ideomotor apraxia	Translating an understood command into a correctly selected and oriented gesture	Left parietal lobe	Compare command, imitation, spontaneous action and object use
Ideational apraxia	Conceptual organisation of an intended action	Contested localisation	Describe the conceptual failure and acknowledge the source conflict
Orobuccal apraxia	Performance of learned oral movements	Left inferior frontal lobe	Separate voluntary task failure from elementary weakness and poor comprehension

TRAP **Do not give a falsely settled localisation for ideational or conceptual apraxia.** The temporal account places it in the left temporal lobe, while the competing account places the stored movement formulae in the left inferior parietal lobule; their only agreement is the left, dominant hemisphere. State the conflict explicitly rather than silently choosing a site.

24.6 Ideomotor apraxia: command, imitation and spontaneous action

The hallmark of ideomotor apraxia is a dissociation: the patient cannot carry out an act to command but may perform the same act spontaneously. The failure concerns selection, sequencing or spatial orientation of the motor act rather than loss of the intention itself. Meaningful as well as meaningless gestures can be affected, as can demonstration of how an imagined tool is used.

Bedside examination should exploit that dissociation. Give a clear verbal command, ask for imitation of a demonstrated gesture, observe spontaneous use, and, where appropriate, compare imagined with real-object use. Imitation and provision of the real object may improve performance. Improvement is informative because it reduces the need to generate the action from a verbal instruction alone. Describe the error: wrong sequence, wrong orientation, an approximation, use of a body part as though it were the object, or inability to initiate a recognisable gesture. Description is safer than merely marking pass or fail.

In right-handed patients, the associated lesions are in the left hemisphere, particularly inferior parietal and prefrontal regions. This fits the broader principle that learned action systems can be dominant-hemisphere functions even when the tested limb itself is not weak. The finding should be interpreted alongside language, sensation and motor examination, because verbal-command testing becomes ambiguous when comprehension is impaired.

A particularly localising dissociation occurs with disconnection. An anterior callosal lesion may leave one limb, usually the left, unable to perform on command while the other limb performs normally. This asymmetric command failure is not ordinary bilateral clumsiness. It points to failure of information transfer to the motor system controlling the affected limb. Alien hand and the wider disconnection syndromes are taught in the dedicated section on alien hand, disconnection, occipital and visual association syndromes.

24.7 Agnosia and cortical sensory signs

In this parietal context, agnosia is best approached through touch. Elementary sensation must reach the cortex, but the patient fails to convert that information into a higher percept or an identified object. This is why the examiner initially establishes that the stimulus is felt and only then asks what its shape, size, traced form or identity means.

Astereognosis is failure to identify a palpated object because perception of its size and shape is impaired. A postcentral-gyrus lesion produces this deficit contralateral to the lesion and may also impair two-point discrimination and graphaesthesia. The patient's inability to name the object is not sufficient by itself. Vision should not be used to solve the task, elementary sensation must be adequate, and a naming disorder must not be mistaken for failed tactile recognition.

Agraphaesthesia refers to inability to recognise a letter or digit traced on the skin. It samples the conversion of a felt trajectory into a learned symbolic form. Compare sides under similar conditions and document whether the patient perceived contact but could not recognise the traced symbol. This distinction turns a vague "sensory problem" into a cortical sensory sign.

Persistent tactile object-recognition failure, or tactile agnosia, follows lesions of the secondary sensory area. The distinction from astereognosis is conceptual: impaired analysis of size and shape prevents recognition in the latter, whereas persistent failure of tactile object recognition defines the former in the permitted account. In either case, the deficit is contralateral to the relevant parietal sensory lesion.

Visual agnosias and other visual association syndromes belong to the dedicated section on alien hand, disconnection, occipital and visual association syndromes. Keeping that boundary matters: "agnosia" names a family of recognition failures, while this section owns the tactile and body-representation signs that localise parietal dysfunction.

24.8 Integrated bedside synthesis

The final answer should sound like a localisation argument, not a catalogue. Begin with the domain that failed, then state the decisive exclusion and the lateralising pattern. A calculation, writing, finger and right-left cluster supports the dominant inferior parietal region. A consistent left-sided attentional bias supports right non-dominant disease. Failure of learned action despite preserved prerequisites establishes apraxia, with the command-spontaneous dissociation sharpening ideomotor apraxia. Contralateral failure of higher tactile recognition supports parietal sensory cortex.

A compact examination sequence is:

- Confirm arousal, comprehension, cooperation, primary sensation, power, coordination and adequate vision for the selected task.
- Test the complete Gerstmann cluster rather than stopping after an isolated writing or calculation error.
- Seek neglect across spontaneous behaviour and sensory modes, then distinguish it from elementary loss and from anosognosia.
- Compare action to command, imitation, spontaneous action and object use; finish with cortical sensory testing on each side.

The formulation should preserve uncertainty. Say “incomplete Gerstmann pattern” when the supported elements cluster but a component is absent; say “possible neglect” when an elementary field or sensory deficit remains unresolved; and say “apraxia cannot be assessed reliably” when comprehension or weakness invalidates the task. These are not evasions. They prevent an invalid test from acquiring false anatomical precision.

Documentation should make the inference auditable. Record the instruction, the patient’s response, the prerequisite that was checked, and the interpretation. “Failed gesture” is weak documentation; “understood the command, had adequate power and coordination, produced a spatially disorganised gesture, then improved with imitation” shows why an ideomotor formulation was considered. Likewise, “ignored the left” is strengthened by noting that elementary input was detected but did not consistently guide behaviour. This style separates observation from localisation and allows another clinician to decide whether the same evidence supports the conclusion.

When findings disagree, resist tidying them into a perfect syndrome. Retest the questionable domain with a simpler instruction, look for a competing prerequisite failure, and describe the residual mismatch. Focal cortical syndromes are valuable precisely because their components can dissociate. Honest discordance is more informative than an overconfident label built from ambiguous tasks.

Finally, integrate associated signs without allowing them to replace the syndrome. Contralateral weakness, hemisensory loss and homonymous hemianopia may follow a lesion in either hemisphere. Constructional apraxia and neglect favour the non-dominant side, while the Gerstmann cluster favours the dominant side. The general ordered bedside cognitive examination is taught in the dedicated section on bedside cognitive examination, executive testing and frontal release signs. Here, the examiner’s task is narrower: use a small set of internally coherent higher cortical signs to identify the parietal operation and hemisphere involved.

Dominant inferior parietal disease produces the Gerstmann pattern; right non-dominant parietal disease produces left hemispatial neglect; apraxia requires preserved elementary function before it can localise, as do cortical sensory signs.

25 · Alien hand, disconnection, occipital and visual association syndromes

Why this cluster belongs together. These syndromes show that intact elementary function is not the same as integrated function. A limb may move and feel yet escape experienced control; vision may be present at the sensory level yet fail to yield recognition; or an entire visual scene may collapse into one attended element. The examination question is therefore not simply whether the patient can move, see or speak. It is whether specialised cortical systems can exchange information and convert perception into ownership, recognition and visually guided action.

The marks lie in a disciplined sequence: define the failed operation, demonstrate it without supplying the answer, identify what remains intact, and localise the interrupted network. Autonomous movement, denial of blindness and recognition failure can otherwise acquire vague psychiatric labels. Bedside analysis replaces them with a testable clinico-anatomical formulation.

RULE **Localise the broken link, not just the visible failure.** In a disconnection or visual-association syndrome, explicitly state the preserved input or output, the higher operation that fails, and the pathway or association cortex whose interruption explains the mismatch.

25.1 Disconnection as a clinical principle

A **disconnection syndrome** arises when an injury severs communication pathways between cortical centres. The relevant centres may each retain some capacity, but information available to one cannot reach the system that must interpret it or act upon it. This produces a characteristic dissociation: the patient succeeds when a task uses an intact route and fails when the same task depends on the disconnected route. That internal contrast is more localising than a general statement that cognition is impaired.

The recognised family includes alexia without agraphia, conduction aphasia, and buccofacial or limb forms of ideomotor apraxia. A comparable connectional principle operates outside the cerebral hemispheres in internuclear ophthalmoplegia from a medial longitudinal fasciculus syndrome. These examples should not be flattened into one mechanism or re-taught as though their bedside signs were interchangeable. Their shared lesson is that a pathway lesion can isolate otherwise functional systems.

The bedside method follows that lesson. First establish that the sensory channel required for the task is usable. Then show that the relevant motor response can occur in another context. Finally, vary the route by which the instruction or object reaches the patient. A route-dependent failure argues for disconnection; failure across every route suggests a more primary sensory, motor or broadly cognitive problem. Language comprehension, attention and cooperation must be considered before assigning a focal label.

Neuroanatomical inconsistency is orderly: performance changes with the pathway demanded by the task. Map that pattern before interpreting the behaviour psychologically. The full assessment of parietal apraxia and agnosia is handled in the dedicated parietal-lobe section.

CALLOSAL	BILATERAL OCCIPITAL	DORSAL	VENTRAL
OWNERSHIP AND INTERHEMISPHERIC CONTROL	CORTICAL BLINDNESS WITH DENIAL	SCENE, GAZE AND VISUAL GUIDANCE	COLOUR PROCESSING

25.2 Alien hand: ownership, agency and control

In **alien hand syndrome**, also called alien limb syndrome, the affected hand retains at least rudimentary motor and sensory function, yet its movements are not under the patient's control and its sensations are not appropriately appreciated. The hand may act in a semi-purposeful, autonomous fashion, such as scratching or unbuttoning. The striking point is not simple involuntary movement. The movement appears organised enough to suggest an intention, but the patient does not experience that intention as his or her own.

The definition contains **two** linked misperceptions: the patient experiences the hand as not belonging to the self, and experiences another agent as governing it. This separates alien hand from a purely motor disorder. Abnormal movement alone does not establish alienness; the disturbance of ownership and agency is central. Conversely, an unusual verbal description should not be accepted without observing the limb and establishing the preservation of basic function.

The affected limb is usually the left hand. It may reach toward or explore objects outside apparent volitional control. It may also counteract the other hand, producing **intermanual conflict**: one hand initiates or completes an act while the other reverses or obstructs it. This opposition gives a visible example of hemispheric intentions failing to remain coordinated.

Bedside description should separate observation from interpretation. Record what the hand did, whether the action was stimulus-bound or apparently spontaneous, whether the patient noticed it, and how ownership and agency were described. Examine elementary motor and sensory capacity, ask for purposeful action, and observe bimanual activity. The aim is to show the mismatch between retained capacity and impaired voluntary integration.

> **GOOD TO REMEMBER**
>
> When an apparently purposeful hand behaves “on its own,” ask separately about **ownership**, **agency** and **control**. Alien hand requires more than an involuntary movement; the patient's experience of possession and authorship is part of the syndrome while rudimentary motor and sensory function remains.

25.3 Callosal and medial-frontal localisation

Alien hand can follow injury to the corpus callosum, and callosal damage within a medial-frontal syndrome can leave the left extremities outside volitional control. The medial-frontal and callosal observations are complementary. Medial-frontal systems contribute to initiation and voluntary control, while the callosum allows the hemispheres to coordinate intention, sensory information and action. Damage at their interface can therefore produce autonomous exploration rather than paralysis.

The classic teaching image is “the right hand not knowing what the left hand is doing,” but the useful anatomy runs in the opposite explanatory direction: language-mediated awareness and control cannot be assumed to reach the systems organising the other limb. This is why laterality matters. The syndrome usually affects the left hand, controlled predominantly by the non-dominant hemisphere, when interhemispheric communication is disrupted by callosal or medial-frontal injury.

Callosal alien hand is an integration disorder. The hand is neither mechanically powerless nor simply insensate. Its activity may look purposive, but the action is not incorporated into the patient's plan. Opposition between hands exposes motor programmes no longer reconciled across the hemispheres.

This distinction prevents psychiatric over-reading. A patient's report that another agent controls the hand resembles the language of passivity, but the clinician should first determine whether there is an observed, lateralised limb phenomenon with retained basic function and a coherent neurological pattern. The statement is then interpreted within ownership and motor-control failure, rather than treated in isolation as evidence of a primary disorder of belief.

PITFALL

Do not diagnose alien hand from the phrase “someone controls my hand” alone. Look for an actual autonomous, semi-purposeful limb action, preserved rudimentary function, impaired ownership and agency, and a reproducible lateralised pattern compatible with medial-frontal or callosal damage.

25.4 The callosal command-apraxia model

A particularly clear disconnection model follows occlusion of both anterior cerebral arteries. The resulting injury involves both frontal lobes and the anterior corpus callosum at the genu. The left arm and leg can move spontaneously, yet they fail to follow verbal or written commands. This is a **left-sided limb apraxia** defined by the route through which action is requested.

The anatomical explanation is an information-transfer failure. Language centres in the left hemisphere cannot pass the command to motor centres in the right hemisphere that organise movement of the left limbs. A command delivered through language therefore fails, while spontaneous movement can remain. The contrast is precisely what makes the syndrome more than weakness: the motor apparatus demonstrates capacity, but the linguistically specified action cannot reach it.

At the bedside, compare spontaneous use with action to spoken and written command. Confirm that the patient understood the instruction and that the limb can move. Describe the deficit as left-sided and command-dependent before offering the callosal localisation. Demonstration by the examiner may use a different route and can help clarify whether the failure is specifically language-mediated, but the interpretation must remain anchored to the observed pattern rather than to one failed gesture.

The broader apraxia taxonomy and parietal examination belong in the parietal-lobe syndromes section. The high-yield point here is narrower: anterior callosal injury can disconnect a comprehended command from the hemisphere controlling the contralateral limbs. Preserved spontaneous movement plus failed command performance supplies the bedside logic.

25.5 Cortical blindness with denial: the Anton presentation

Anton syndrome is denial of cortical blindness. The patient behaves as though vision were normal and may deny blindness directly or reveal the denial through conduct. Confabulation is frequent: the patient may describe the room, clothing or objects despite being unable to see them. The syndrome is therefore not merely blindness plus a verbal claim. It is cortical blindness accompanied by unawareness of the deficit, expressed verbally or through behaviour.

The localisation requires bilateral occipital damage in the posterior cerebral artery territories. An illustrative vascular pattern is a new right posterior cerebral artery stroke superimposed on an earlier left-sided stroke, leaving bilateral occipital injury and cortical blindness. The sequence teaches an important principle: a patient can appear visually confident because awareness of loss is itself impaired.

Examination begins with behaviour, not confrontation. Observe navigation, reaching and responses to the environment; then establish visual failure and ask the patient to explain any errors. If the patient supplies confident visual descriptions, record them as confabulatory responses only after confirming that the stated information could not have been obtained visually. The examiner should distinguish a failure to acknowledge the deficit from an ordinary attempt to preserve dignity or from uncertainty during testing.

The immediate consequence is safety. Denial of blindness commonly undermines adherence to ward arrangements, prompts demands for early discharge and can lead to dangerous behaviour. A cortically blind patient acting as though fully sighted may endanger self and others and must not leave unsupervised. Capacity and disposition decisions require the actual functional risk to be documented, not merely the diagnostic label.

25.6 Visual agnosia and recognition of objects

Visual agnosia is inability to comprehend or identify an object that is seen despite an intact visual system. The dissociation lies between seeing features and knowing what those features mean. A patient may describe a stop sign as red, octagonal and bearing the letters S-T-O-P, yet be unable to state what a driver should do. Elementary visual analysis has supplied colour, shape and lettering, but the percept has not acquired usable object meaning.

This distinction guides examination. First ensure that the patient can access the visual features needed for the task. Ask for a description before asking for the object's name or use; this reveals whether shape, colour or component detail is available. Then test whether identification improves through a nonvisual route. The resulting profile should be described rather than reduced prematurely to “poor vision” or “memory loss.” The defining problem is recognition from vision despite an intact visual system.

Visual agnosia is detected most often in association with multiple small strokes or Alzheimer disease. That association does not make either diagnosis from the bedside sign alone. The syndrome identifies the failed cognitive operation; aetiology requires the broader history, examination and investigation. The full dementia syndrome and its diagnostic workup belong to the Dementias and neurocognitive disorders notes.

Severe visual agnosia, termed “psychic blindness,” is a major component of Kluver-Bucy syndrome and can drive oral examination of objects. This is a connection to, not a substitute for, the dedicated temporal-lobe discussion of that syndrome. For the present section, its value is conceptual: when vision no longer yields object identity, the patient may seek information through another sensory route.

25.7 Prosopagnosia: a selective failure of face recognition

Prosopagnosia is inability to recognise familiar faces despite an intact visual system. The deficit is selective rather than a failure to recognise the person in every modality. Patients can identify individuals from voice, dress and mannerisms, while object recognition and reading remain intact. The examiner should therefore ask how recognition occurred instead of accepting a correct name as proof that facial recognition is preserved.

The clean localisation is occipitotemporal dysfunction. Lesions may involve bilateral inferomedial visual-association cortices or the right occipitotemporal lobe. Structural injury and neurodegeneration, including Alzheimer disease and frontotemporal dementia, can produce the syndrome. These alternatives again separate syndrome from cause: the same selective recognition failure can arise in an acute focal setting or during a progressive disorder.

Bedside assessment requires familiar stimuli whose identity can be checked without giving contextual clues. Ask the patient to identify the person from the face, then establish whether voice, dress or mannerisms permit recognition. Also confirm that reading and ordinary object recognition are usable. The characteristic dissociation is not “cannot name a photograph”; it is impaired identification by face with successful identification through other personal cues.

Localisation language must remain precise. Occipitotemporal cortex, bilateral inferomedial visual-association cortex and the right occipitotemporal lobe provide the regional formulation. Do not add a narrower gyral label.

TRAP **Do not collapse all visual-association failure into “visual agnosia.”** Object meaning, familiar-face identity, awareness of blindness, attention to a scene, voluntary gaze, visually guided reaching and colour perception are separable operations with different clinico-anatomical patterns.

25.8 Balint syndrome and the dorsal visual pathway

Balint syndrome is the archetypal disorder of the dorsal visual pathway and follows bilateral parieto-occipital damage, which may arise from strokes or Alzheimer disease. It contains a **three**-part syndrome: simultanagnosia, psychic paralysis of gaze or ocular apraxia, and optic ataxia.

Each component examines a different transformation of visual information, so merely naming the triad earns less than showing how its parts fit together.

Simultanagnosia is inability to grasp more than one visual object at once. When shown a complex scene, the patient may report only a single element, which may even be peripheral. This is not simply poor visual acuity. The failure concerns simultaneous visual attention and integration: isolated parts can enter awareness, but the scene does not become a coherent whole.

Ocular apraxia, also described as psychic paralysis of gaze, is inability to direct gaze voluntarily. **Optic ataxia** is misreaching under visual guidance. Together they show why the dorsal pathway is clinically concerned with more than static perception. It helps organise attention, gaze and action in relation to the visible world. The patient may see an object yet fail to orient the eyes toward it or guide the hand accurately to it.

Examine the components separately. Use a complex scene to assess whether more than one element can be integrated; ask the patient to shift gaze voluntarily toward selected targets; and observe reaching to visible objects. Before interpreting failure, ensure that the patient has understood the task and that gross movement and elementary vision are adequate for it. The triad is strongest when its parts converge on the same bilateral parieto-occipital, dorsal-stream account.

25.9 Ventral colour processing and an integrated bedside map

Colour is processed through the ventral visual stream. The human colour-perception area lies along the collateral sulcus and fusiform gyrus. **Cerebral achromatopsia** follows bilateral or non-dominant inferior occipitotemporal lesions involving the lingual and fusiform gyri while sparing the calcarine cortex. The pattern differs from cortical blindness because primary visual cortex is spared, and from prosopagnosia because the failed visual attribute is colour rather than familiar-face identity.

The examination should isolate colour perception from language and object recognition as far as the clinical setting allows. Establish that the patient can see the material, then ask for comparison, matching or description of colours without relying on a single naming response. The purpose is to determine whether colour itself is unavailable or whether the patient sees it but cannot retrieve a verbal label. This is a clinical reasoning sequence, not a substitute for formal visual assessment.

SYNDROME	FAILED OPERATION	PRESERVED CLUE	SUPPORTED ANATOMY OR ASSOCIATION
Alien hand	Ownership, agency and voluntary integration of limb action	Rudimentary movement and sensation	Medial frontal or corpus callosum
Anton syndrome	Awareness of cortical blindness	Confident behaviour may conceal the loss	Bilateral occipital cortex
Visual agnosia	Meaning of a seen object	Elementary visual features can be described	Most often detected with multiple small strokes or Alzheimer disease
Prosopagnosia	Identity from a familiar face	Voice, dress, mannerisms, reading and object recognition	Bilateral inferomedial or right occipitotemporal cortex
Balint syndrome	Scene integration, voluntary gaze and visually guided reaching	Individual capacities may be demonstrable when isolated	Bilateral parieto-occipital dorsal pathway
Cerebral achromatopsia	Colour perception	Calcarine cortex is spared	Inferior occipitotemporal lingual and fusiform regions

MNEMONIC

SEE - KNOW - FACE - SCENE - COLOUR

See: establish cortical vision and awareness. **Know:** test object meaning. **Face:** test familiar-face identity without voice or dress cues. **Scene:** test simultaneous attention, gaze and visually guided reach. **Colour:** test colour perception separately. The sequence keeps neighbouring syndromes from being collapsed into one vague “visual problem.”

Build the integrated answer from dissociations. Ask what the patient can detect, what meaning can be extracted, whether awareness of loss is present, and whether vision can guide gaze and reach. Test each operation separately, then seek a coherent pattern of preserved and impaired functions.

- Describe the behaviour before naming the syndrome.
- Establish elementary sensory and motor capacity required by the task.
- Change the input route or response demand to expose a disconnection.
- State the failed operation, preserved function and supported localisation in one formulation.

Disconnection syndromes are diagnosed by dissociation: demonstrate what remains intact, identify the higher operation that fails, and localise the pathway that should have joined them.

26 · Pseudobulbar affect: phenomenology, pathophysiology and treatment

Why this syndrome earns bedside marks. **Pseudobulbar affect** is a disorder of emotional expression in which laughing or crying escapes voluntary regulation and no longer maps reliably onto the patient's prevailing mood. The same construct is described as pathological laughing and crying, involuntary emotional expression disorder, emotional incontinence or emotional lability. Its dramatic appearance invites a psychiatric shortcut, especially when crying is mistaken for depression. The better answer separates expressed emotion from experienced emotion, demonstrates the associated neurological syndrome, and then explains the disrupted motor-affective network.

> **IN INDIA**
>
> For pseudobulbar affect, options include an SSRI or a tricyclic such as amitriptyline; no comparative hierarchy is established, so individualise the choice. Explain the treatment to the family as reducing the outbursts rather than as treating low mood.

The topic joins phenomenology, examination and mechanism. Phenomenology establishes that the episode is brief, stereotyped, excessive and uncontrollable. Examination asks whether corticobulbar upper motor neuron signs or a broader cerebral disorder are present. Pathophysiology then moves from the traditional release model to a distributed fronto-ponto-cerebellar account, while treatment targets the expression disorder without assuming that the patient's mood has changed. This sequence is clinically useful because pseudobulbar affect and depressive illness can coexist.

RULE **Diagnose a mismatch between emotional display and underlying mood, not crying alone.** Establish the abrupt, uncontrollable and disproportionate episode; ask what the patient actually felt before, during and after it; and examine for the neurological disorder that makes impaired emotional regulation plausible without requiring residual sadness or happiness.

26.1 The emotional burst: core phenomenology

The conspicuous clinical feature is **emotional lability**. An episode begins suddenly, usually as crying and less often as laughter, and is excessive in relation to the conversation or situation. It may be triggered by something emotionally relevant, but the magnitude and persistence of the display exceed what the patient feels. At other times the outburst appears unprovoked. The defining abnormality is therefore not the presence or absence of a trigger; it is loss of proportional, voluntary control over expression with an increased frequency of laughing or crying.

Episodes are stereotyped for the individual and last **from a few seconds to a few minutes**. The patient cannot reliably suppress them once they begin, yet often recognises that the display is inconsistent with, or excessive relative to, the underlying feeling. When the burst ends, there is no persisting sadness after crying and no persisting happiness after laughter. That rapid return to the premorbid emotional state is among the most useful bedside clues.

Interview the event as an event. Ask for its onset, observable form, trigger, controllability, duration, stereotypy, frequency and emotional aftermath. Then obtain an informant account when available, because embarrassment, poor recall or adaptation may lead patients to underdescribe visible episodes. Record the patient's own appraisal: “I was not that sad,” “I could not stop,” or “the laughter did not fit” is more informative than a global report of being emotional. Avoid leading questions that collapse felt mood and motor expression into the same construct.

MNEMONIC

BURST

Brief; **U**ncontrollable; **R**ecurrent or unusually frequent; **S**tereotyped and stronger than the felt emotion; **T**erminates without a residual corresponding mood. Each element is supported by the characteristic episode phenomenology of pathological laughing and crying.

26.2 Clinical diagnosis and the depression differential

There are no established formal diagnostic criteria for pseudobulbar affect. Diagnosis remains clinical: identify the characteristic expressive episodes, establish their relationship to subjective mood, and interpret them in the context of neurological disease. Rating scales can quantify symptoms, but a score cannot replace the interview or the observation of an episode. Among available measures, the **Pathological Laughter and Crying Scale** was developed specifically for pathological emotion after stroke. Its proper role is structured severity assessment rather than conversion of a heterogeneous phenomenon into a stand-alone criterion set.

The main psychiatric distinction is from depression after stroke or brain injury. Pseudobulbar crying is abrupt, brief, stereotyped, difficult to control and out of proportion to the trigger. It may be incongruent with the prevailing mood and stops without residual sadness. Depression is organised around sustained low mood, with congruent affect and accompanying neurovegetative symptoms. The same logic separates pseudobulbar laughter from a sustained elevation of mood: the visible act alone is not the mood state.

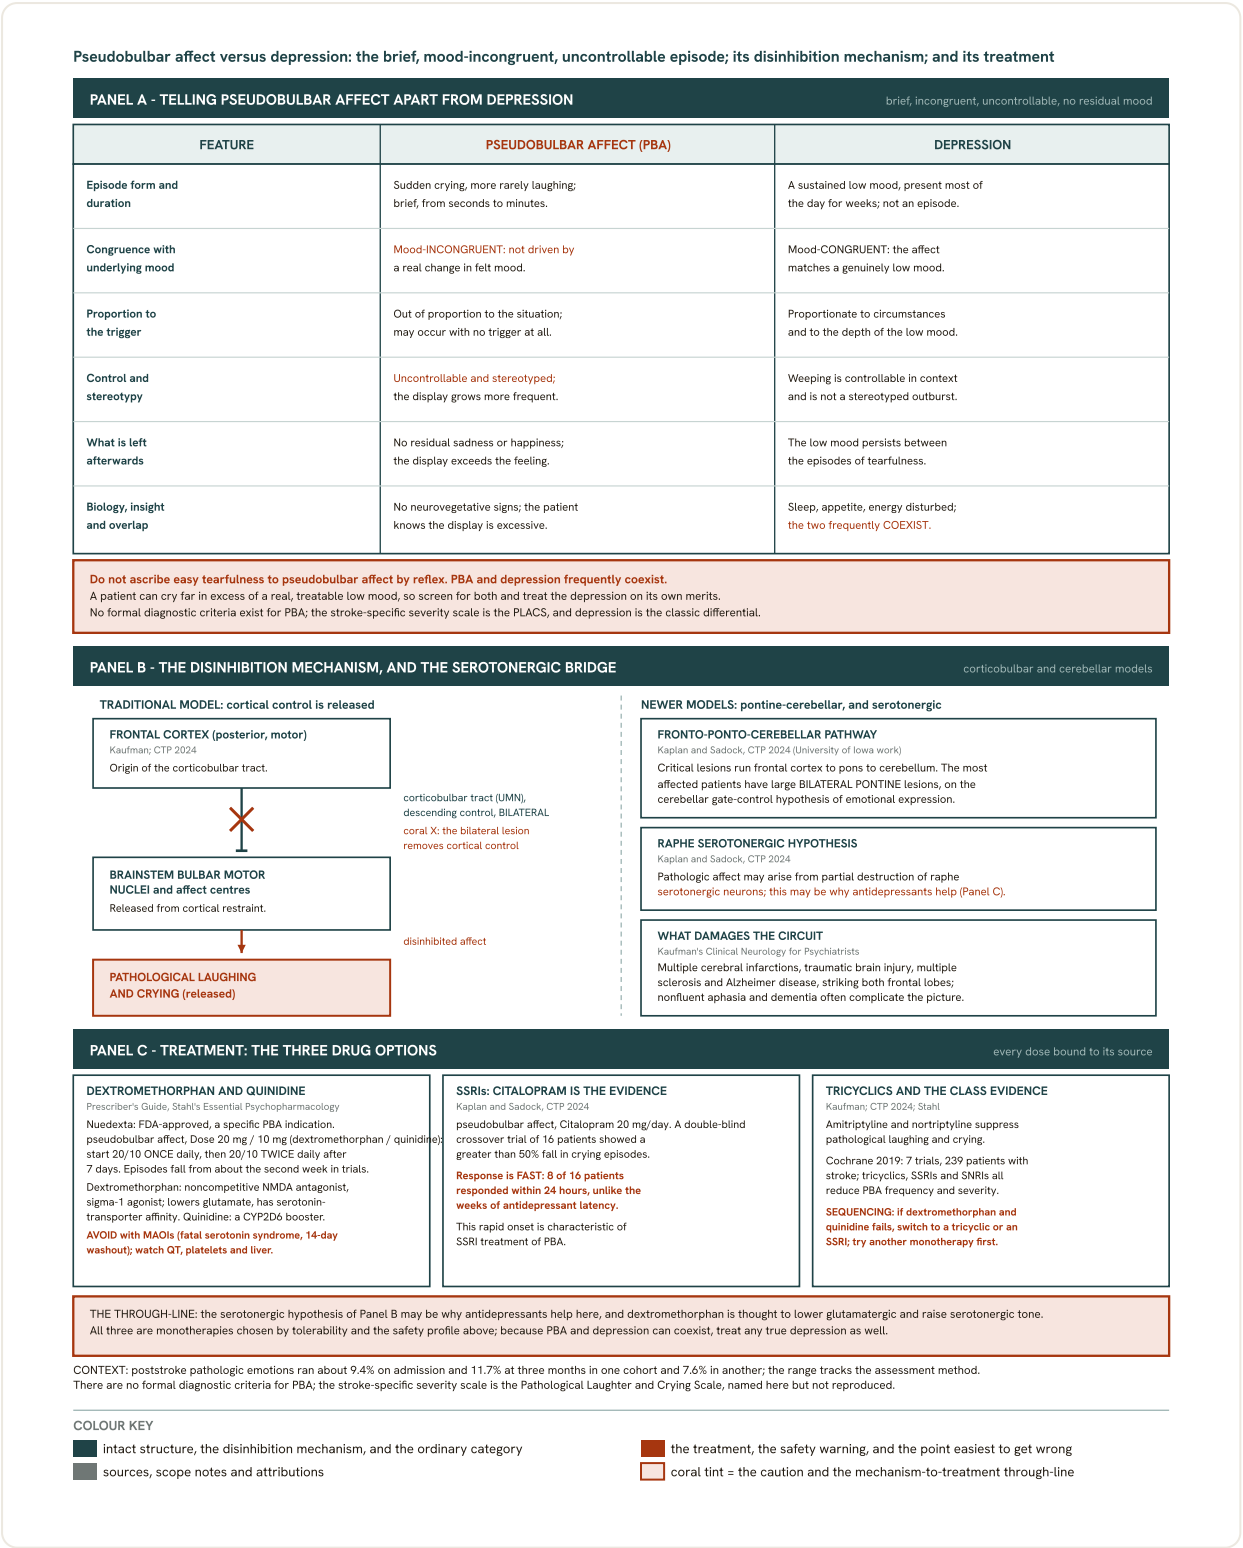


Fig 34.9 Pseudobulbar affect versus depression: the episode, the mechanism, and the treatment. Pseudobulbar affect is a brief, stereotyped, mood-incongruent and uncontrollable outburst of crying or laughing that is out of proportion to its trigger and leaves no residual feeling, unlike the sustained, congruent low mood of depression, though the two frequently coexist. It is traditionally read as bilateral corticobulbar lesions releasing brainstem affect centres from cortical control, with newer fronto-ponto-cerebellar and raphe-serotonergic models; the serotonergic account may explain why SSRIs and tricyclics help, while dextromethorphan acts mainly through NMDA and sigma-1 mechanisms, with quinidine added as a pharmacokinetic booster. (Kaufman's Clinical Neurology for Psychiatrists; Kaplan and Sadock, CTP 2024; Prescriber's Guide, Stahl's Essential Psychopharmacology.)

The accompanying figure reinforces the crucial separation between episode and syndrome. The following table turns it into an answer that can be used at the bedside.

AXIS	PSEUDOBULBAR AFFECT	DEPRESSIVE ILLNESS
Temporal pattern	Sudden, brief, stereotyped outbursts	Sustained low mood
Relation to trigger	Disproportionate or apparently unprovoked	Affect is generally understandable within the mood syndrome
Display versus feeling	Display exceeds, or is incongruent with, the felt emotion	Affect is congruent with the low mood
After the episode	No residual sadness or happiness attributable to the burst	Low mood persists beyond an individual crying spell
Associated symptoms	Neurological context and impaired control of emotional expression	Neurovegetative features support the mood syndrome
Relationship between them	Pseudobulbar affect and depression may coexist; demonstrating one does not exclude the other.	

PITFALL

Do not re-label every easy tear as pseudobulbar affect in a patient with a neurological lesion. Ask about sustained mood and neurovegetative change as well as the episode itself. Conversely, do not diagnose depression merely because the patient cries. The expression disorder and depressive illness frequently coexist and require separate formulations.

26.3 Neurological setting and the bulbar examination

Pseudobulbar affect can accompany disorders that injure frontal systems, descending corticobulbar pathways or more extensive cerebral networks. Supported settings include multiple cerebral infarctions, traumatic brain injury, multiple sclerosis and Alzheimer disease. When both frontal lobes or a wider area of cerebrum is affected, nonfluent aphasia and dementia may complicate assessment. The examiner may then need collateral description and direct observation because the patient cannot give a detailed account of internal mood or loss of control.

Pseudobulbar palsy reflects upper motor neuron corticobulbar disease. Pseudobulbar and **bulbar palsy** share dysarthria, dysphagia and absent voluntary palate movement, but their accompanying signs differ. The distinction matters because emotional lability belongs to the pseudobulbar pattern, whereas bulbar palsy is a lower motor neuron syndrome and does not explain pathological emotional expression.

BEDSIDE FEATURE	PSEUDOBULBAR PALSY	BULBAR PALSY
Motor level	Upper motor neuron, corticobulbar	Lower motor neuron
Shared bulbar difficulty	Dysarthria, dysphagia and loss of voluntary palate movement	
Gag reflex and jaw jerk	Hyperactive	Hypoactive
Respiratory impairment	Absent	Present
Intellectual impairment	May be present	Absent in the defining comparison
Emotional lability	Present	Absent

Post-stroke frequency estimates vary with timing and who supplies the history. This is not statistical clutter; it explains why a patient-only interview may miss a public, episodic behaviour.

9.4%	11.7%	17.9%	6.3%
DURING STROKE ADMISSION	SAME COHORT, THREE-MONTH FOLLOW-UP	CAREGIVER-ASSISTED, THREE-MONTH ASSESSMENT	PATIENT-ONLY ESTIMATE IN THAT STUDY

A consecutive cohort of 508 people with ischaemic stroke generated the admission and later estimates, while a separate assessment of 127 people showed the caregiver-versus-patient contrast. Another cohort of 516 people assessed three months after stroke yielded an estimate of 7.6%. These figures should remain explicitly post-stroke; the evidence supplied here does not justify exporting them to other neurological diseases.

26.4 From corticobulbar release to a distributed circuit

The traditional explanation is **corticobulbar disinhibition**. The corticobulbar tract originates in the posterior frontal motor cortex and descends to brainstem motor nuclei, which then innervate the muscles of the head and neck. In the classical model, bilateral interruption of descending neocortical upper motor neuron control releases brainstem affective motor centres. Emotional expression can then be triggered too readily and cannot be shaped or stopped by cortical intention in the same way that bilateral corticobulbar injury produces pseudobulbar palsy.

That account is valuable but incomplete. A major clinical text states that the bilateral lesion pattern is not uniform: frontal or temporal lesions in either hemisphere are also associated with pathological emotion. The observed syndrome therefore cannot be reduced to a single tract cut at a fixed point. It is better viewed as loss of regulation within a network that estimates the emotional context, selects an appropriate motor response, scales it and corrects it while it unfolds.

TRAP **"Pseudobulbar affect requires bilateral corticobulbar lesions" is the textbook rule, not a universal clinical law.** The classic pseudobulbar-palsy account specifies bilateral corticobulbar tract damage, but clinical observations report pathological emotion with unilateral frontal or temporal lesions in either hemisphere. State the rule and the exception. Do not erase the conflict or resolve it by assertion.

The broader model places the critical disturbance along the **fronto-ponto-cerebellar pathway**. Frontal systems provide contextual and intentional signals; pontine relays carry them toward the cerebellum; and the cerebellar contribution can be understood as modulation of the intensity and timing of the emotional motor response. In a lesion analysis of 12 patients with pathological crying, the heaviest episode burden occurred with relatively large bilateral pontine lesions. This supports a gate-control formulation without making one lesion site obligatory.

A complementary **serotonergic hypothesis** proposes partial destruction of raphe serotonergic neurons or their projections, particularly in relation to pontine damage. It provides a mechanistic bridge between lesion anatomy and the rapid benefit observed with serotonergic antidepressants. The hypothesis should be presented as explanatory support, not as a proven biomarker or a complete account of every case.

26.5 Management frame: locus, focus and modus

Treat the target that is actually present. LOCUS depends on medical and neurological stability. A patient with an established, stable neurological disorder and characteristic recurrent episodes can usually be assessed and followed in outpatient care. A new presentation with an evolving focal deficit, altered consciousness, swallowing risk or diagnostic uncertainty requires urgent neurological and medical assessment; suppressing the visible crying is not a substitute for identifying a new lesion or systemic problem.

FOCUS begins with the acute burden: loss of control, episode frequency and severity, embarrassment, interrupted communication and risk during swallowing, mobility or care. In the short term, establish whether the chosen treatment reduces episodes without producing disproportionate adverse effects. Longer-term review separates residual pseudobulbar episodes from depression, cognitive decline and the underlying neurological disease. Improvement in crying does not prove that a comorbid depressive syndrome has remitted.

MODUS combines explanation, environmental adaptation and pharmacotherapy. Explain to the patient and family that the expression is involuntary and may exceed the actual feeling. This reduces moral interpretations such as weakness, rudeness or emotional manipulation. Agree on a simple way to pause conversation during an episode, allow recovery, and resume without repeatedly demanding an explanation. Address the consequences in rehabilitation, family interaction and public settings while continuing care of the underlying neurological disorder.

- Define the treatment target in observable terms: frequency, severity, controllability and disruption caused by episodes.
- Assess sustained mood independently, because depression and pseudobulbar affect may coexist.
- Choose a supported monotherapy, review benefit and tolerability, and avoid confusing rapid control of expression with an antidepressant response.
- Recheck diagnosis, adherence, interactions and the neurological course before declaring treatment failure.

Pharmacological options supported in the supplied evidence are dextromethorphan with quinidine, selective serotonin reuptake inhibitors and tricyclic antidepressants; selective serotonin-norepinephrine reuptake inhibitors also have class-level trial support after stroke for reducing the frequency and severity of pathological affect. Use current product information and local prescribing guidance for screening and monitoring because the combination treatment has clinically important cardiac, hepatic, haematological and interaction risks.

26.6 Dextromethorphan with quinidine

Dextromethorphan with quinidine is FDA-approved specifically for pseudobulbar affect under the trade name Nuedexta. Dextromethorphan is a noncompetitive NMDA receptor antagonist and sigma-1 receptor agonist; it reduces glutamatergic neurotransmission and also has affinity for the serotonin transporter. These actions provide several plausible routes by which it may alter the dysregulated emotional-expression circuit without implying that pseudobulbar affect is simply a glutamate or serotonin disorder.

Quinidine is not included here to treat an arrhythmia. It acts as a **pharmacokinetic booster** by inhibiting CYP2D6 metabolism of dextromethorphan, thereby increasing dextromethorphan exposure. This distinction predicts efficacy variation as well as interaction risk. People who are poor CYP2D6 metabolisers may gain little additional dextromethorphan availability while remaining exposed to quinidine adverse effects.

For pseudobulbar affect, begin dextromethorphan 20 mg with quinidine 10 mg once daily. **After 7 days, for pseudobulbar affect, increase to dextromethorphan 20 mg with quinidine 10 mg twice daily**, which is the usual regimen. Some patients may tolerate and respond to doses above the approved regimen, but controlled evidence for such dosing is limited, and the supplied source gives no explicit numerical maximum for pseudobulbar affect.

Set expectations around the target rather than promising mood change. In clinical trials, episode frequency fell significantly beginning on **day fifteen**, with benefit expressed as reduced frequency and severity of uncontrollable laughing or crying. The permitted evidence does not provide a placebo-adjusted effect size, so a precise percentage reduction should not be quoted. Review an episode record and collateral observations alongside tolerability.

26.7 Antidepressants and treatment sequencing

Antidepressants can suppress pathological laughing and crying even when the clinical target is not a depressive disorder. The serotonergic lesion hypothesis offers one rationale, while the characteristically rapid response in pseudobulbar affect helps distinguish control of emotional expression from the slower change expected when treating a sustained mood syndrome. This does not make antidepressants diagnostic tests; response remains supportive and must be interpreted with the original phenomenology.

A Cochrane Review published in 2019 included seven trials and 239 people with stroke. Across the small evidence base, tricyclic antidepressants, selective serotonin reuptake inhibitors and selective serotonin-norepinephrine reuptake inhibitors significantly reduced the frequency and

severity of pathological affect. The finding supports a class-level treatment option after stroke, but does not erase differences in adverse effects, interactions or patient vulnerability.

The most concrete selective serotonin reuptake inhibitor example is post-stroke pathological crying treated with citalopram. In a double-blind, placebo-controlled crossover study of 16 patients, crying episodes decreased by more than 50%. For pseudobulbar affect, citalopram 20 mg per day produced a response **within 24 hours in eight patients**; three responded within 3 days and four required more than a week. The rapid onset is the memorable point, not a promise that every patient will respond immediately.

Tricyclic antidepressants are also established options. Amitriptyline is reported to suppress pathological laughing and crying, and an illustrative post-stroke case improved greatly with nortriptyline. No treatment dose for either tricyclic is provided in the permitted evidence, so none should be inferred from general antidepressant prescribing.

GOOD TO REMEMBER

If dextromethorphan with quinidine is ineffective, consider switching to a tricyclic antidepressant or a selective serotonin reuptake inhibitor. With a partial response, try another appropriate monotherapy before moving to augmentation as the preferred sequencing principle. Reassess whether the residual problem is pseudobulbar affect, depression or their coexistence.

26.8 Safety, follow-up and the integrated answer

Dextromethorphan with quinidine demands an interaction and cardiac screen before prescribing. It must not be combined with a monoamine oxidase inhibitor because the combination can cause fatal serotonin syndrome, and it must not be started for **at least 14 days after stopping a monoamine oxidase inhibitor**. It is also contraindicated in complete atrioventricular block without a pacemaker, or when the risk of complete block is high, and with a QT-prolonging drug metabolised by CYP2D6, such as thioridazine or pimozide.

CYP2D6 inhibition is clinically relevant beyond dextromethorphan: quinidine can raise concentrations of other CYP2D6 substrates, with desipramine as a supported example. Medication reconciliation should therefore include prescribed, over-the-counter and recently stopped drugs. The interaction review is not a paperwork exercise: it is part of deciding whether this otherwise targeted treatment is suitable for the individual patient.

Dangerous adverse effects include immune-mediated thrombocytopenia, hepatotoxicity and dose-dependent QT prolongation. Thrombocytopenia may be severe or fatal; if it occurs, discontinue immediately unless it is clearly unrelated to the drug, and do not restart in a sensitised patient. Quinidine is also associated with a lupus-like syndrome with polyarthritis. These risks explain why dose escalation beyond the supported regimen should not be improvised merely because episodes persist.

Follow-up should return to the original formulation. Ask whether bursts are less frequent, less severe, more controllable or less disruptive; obtain collateral observation when feasible; and repeat the mood assessment separately. If there is no benefit, confirm exposure, reconsider the

diagnosis and choose another supported monotherapy rather than stacking drugs reflexively. If episodes improve but sustained sadness remains, treat the depressive syndrome on its own evidence base; general antidepressant pharmacology and comparative drug selection are covered in the Mood disorders: pharmacotherapy notes.

An integrated examination answer is concise but layered: define the mismatch between display and feeling; describe the brief, stereotyped and uncontrollable episode; distinguish it from sustained mood change; examine the corticobulbar syndrome; state the classical bilateral release model and its unilateral-lesion exception; then outline supported monotherapies and the quinidine-driven safety checks. That structure turns an easily recognised symptom into a defensible neuropsychiatric formulation.

Pseudobulbar affect is a brief, stereotyped and uncontrollable disorder of emotional expression: the display exceeds the felt mood, bilaterality is the classic lesion rule but not universal, and treatment must target episodes while depression is assessed separately.

Structural lesions and the seizure overlap

27 · Chronic subdural haematoma as reversible decline in the elderly

A declining older adult may be carrying a removable lesion, not an irreversible diagnosis. **Chronic subdural haematoma** earns marks because it joins neurology, old-age psychiatry and emergency judgement. It can present through headache, confusion, altered personality and cognitive impairment while producing surprisingly little focal neurology. The clinician who waits for a dramatic hemiparesis, a witnessed major injury or a tidy neurological history may therefore miss it.

The central examination proposition is simple: this is a structural mimic of psychiatric illness and progressive cognitive disorder whose timely recognition can restore cerebral function. Structural imaging is not an ornamental investigation in such a case. It is the step that converts a vague syndrome into a potentially correctable diagnosis. The full dementia syndrome and its general diagnostic work-up belong in *the Dementias and neurocognitive disorders notes*; the present section teaches the subdural lesion that must not be overlooked within that work-up.

Venous	Insidious	Subtle	Treatable
ORIGIN	TEMPO	FOCAL SIGNS	PRINCIPLE

27.1 What the lesion is

The **subdural space** lies between dura and arachnoid. A subdural haematoma is therefore **extra-axial**: blood collects outside the brain rather than within cerebral tissue. The usual source is torn **bridging veins**. Because these veins bleed at relatively low pressure, blood may ooze through an extensive potential space before the expanding collection meets sufficient resistance from the underlying brain. The collection is commonly spread over a cerebral hemisphere and may be bilateral.

This vascular origin explains the tempo. A low-pressure venous collection can enlarge without the immediately catastrophic presentation expected from high-pressure arterial bleeding. Epidural haemorrhage is classically arterial, often involving the middle meningeal artery, and subarachnoid haemorrhage usually follows arterial bleeding from a ruptured aneurysm. The contrasts matter because “intracranial bleed” is not a single physiological event. Site, vessel and pressure determine how quickly symptoms appear and which clinical signals dominate.

Chronicity also changes what the clinician sees. A collection that develops and persists over weeks may spread widely in the subdural compartment. Generalised cerebral dysfunction then outweighs clean lateralising signs. Headache, behaviour and cognition may carry more diagnostic weight than the expectation of a dense motor deficit.

MNEMONIC

SLOW

Subtle focal signs; **L**ucent aged blood on CT; **O**lder person with falls, bleeding-risk medicines or atrophy; **W**orsening headache, personality and cognition. Each letter points back to the lesion's quiet venous expansion rather than to a separate syndrome.

27.2 Why the elderly are vulnerable

Susceptibility rises in people **older than 65 years** because several risks converge. Falls increase the opportunity for head movement and venous tearing. Aspirin, anticoagulants and other medicines that increase bleeding tendency make a small vascular injury harder to contain. Most importantly, **age-related cerebral atrophy** places greater tension on the bridging veins and reduces the brain's capacity to compress a bleeding vessel once bleeding has begun.

That mechanism corrects a common misconception: the severity of the initiating trauma need not match the later disability. In an atrophied brain, even minimal trauma may tear delicate meningeal bridging veins. A remembered fall may be absent, dismissed as trivial or never observed. The history “no head injury” therefore lowers neither suspicion nor urgency enough to exclude the lesion. The relevant question is not only whether a dramatic injury occurred, but whether the biological setting permits a small injury to continue bleeding.

Medication history must be exact in purpose even when the psychiatric presentation appears dominant. Ask what is taken, why it is taken, who supervises it and whether recent dosing has changed. This is not an instruction to stop essential antithrombotic treatment blindly. It is a prompt to recognise bleeding risk and coordinate urgent decisions with the treating medical and neurosurgical teams. The same history should include falls that relatives regard as inconsequential.

GOOD TO REMEMBER

A reliable collateral history is often more diagnostic than the patient's account. Ask separately about minor falls, new headache, functional slippage, altered social conduct, reduced self-care and all medicines that impede clotting. If the story is inconsistent with the injuries or the care environment seems unsafe, consider elder abuse as a possible source of head trauma and activate safeguarding procedures.

27.3 The psychiatric and cognitive masquerade

The usual syndrome is insidious rather than spectacular. Headache is most often the first symptom of chronic subdural haematoma. As cerebral dysfunction develops, confusion, personality change and cognitive impairment may emerge. These are nonspecific manifestations, which is precisely why the patient may first reach psychiatry, a memory clinic or a general physician rather than a neurosurgical service.

The diagnostic clue is the mismatch between broad behavioural or cognitive change and only subtle focal physical findings. Families may describe altered conduct, reduced organisation or inability to manage familiar routines. Such descriptions should not be prematurely translated into a primary psychiatric or degenerative disorder. They identify a change in brain function; they do not identify its cause. The clinician must preserve the structural differential until imaging has addressed it.

The syndrome can resemble dementia because decline is sustained and functionally important, yet it can also be confused with delirium when confusion dominates. Those syndromic diagnoses have their own criteria and examinations elsewhere in this chapter and in the dedicated neurocognitive notes. Here, the useful move is aetiological: ask whether an extra-axial collection can explain the observed decline. Do not allow a psychiatric label to become a reason not to image.

Diagnostic closure often occurs through an apparently coherent narrative. An older person becomes forgetful, stops managing money, neglects household tasks and appears emotionally altered; the family then supplies the word “dementia”, and every later observation is interpreted through it. Chronic subdural haematoma breaks that narrative because one lesion can generate the whole cluster. The better formulation is provisional and layered: describe the cognitive and behavioural syndrome, identify its tempo, record the headache and risk context, and keep the structural cause open until imaging has answered it.

The reverse error is also possible. A patient may already have cognitive impairment or a psychiatric disorder, and new deterioration is attributed to the established diagnosis. Baseline vulnerability does not protect against a superimposed subdural collection. The question is whether there has been a meaningful change from that person's own baseline. A recent shift in conduct, self-care, communication or mobility deserves fresh causal assessment even when the case file already contains a persuasive label.

TRAP **No major trauma and no obvious lateralising deficit do not exclude chronic subdural haematoma.** The lesion often follows little or unremembered trauma and produces generalised headache, confusion, personality change and cognitive impairment more prominently than focal signs.

27.4 Bedside diagnostic reasoning

Assessment should reconstruct change, not merely record symptoms. Establish the previous level of cognition, personality and everyday functioning, then ask how the present state departed from it. The temporal sequence matters: headache preceding behavioural and cognitive decline supports a mass-lesion frame more strongly than a timeless list of complaints. A complete neurological examination remains necessary, but a quiet examination does not cancel the history.

A focused bedside enquiry should cover:

- new or evolving headache, including whether it was the earliest change noticed;
- confusion, personality alteration and cognitive or functional decline, preferably corroborated by someone who knew the baseline;
- falls, antiplatelet or anticoagulant exposure, and the context of cerebral atrophy as converging vulnerabilities;
- minor or absent recalled trauma and possible inflicted injury, without using a negative trauma history as a rule-out test.

Several differentials may coexist at the point of presentation: a primary neurocognitive or psychiatric disorder, delirium, medication effects, brain tumour, cerebrovascular disease and hydrocephalus. Listing them is not enough. The psychiatrist's responsibility is to rank the reversible and dangerous possibilities early. A chronic subdural haematoma deserves priority because clinical phenomenology alone cannot safely remove it from consideration and structural imaging can reveal the cause.

Collateral information should be anchored to observable change. “More confused” is less useful than examples of missed meals, unfamiliar errors, altered judgement or a new inability to complete a previously routine activity. Ask whether headache appeared before, with or after the functional decline. Reconstruct falls from witnesses, bruising, broken objects and changes in walking confidence rather than relying on the patient's spontaneous recall. None of these observations diagnoses the lesion, but together they determine whether a bland neurological examination is falsely reassuring.

Examination findings should likewise be interpreted as a pattern. Generalised cerebral dysfunction with only slight asymmetry fits the established presentation; severe focal findings are not required. At the same time, any new focal abnormality increases urgency rather than changing the case into a purely stroke-shaped problem. Record consciousness, behaviour, cognition and the neurological examination in language that can be compared across assessments. Serial change is often more informative than an isolated adjective such as “stable”.

Interpret cognitive test performance as evidence of impairment, not evidence of aetiology. The standardised instruments and their scoring are covered in *the Psychotherapies, clinical assessment, MSE and nosology notes*. In this setting, a low score may document the consequence of cerebral dysfunction, but it cannot distinguish a subdural collection from other structural or medical causes.

27.5 CT: density changes with the age of blood

Structural neuroimaging is the decisive investigation. CT is generally sufficient to screen for subdural haematoma in the dementia work-up. The image must be read in relation to blood age: fresh blood is denser than brain, whereas aged blood in a chronic collection is less dense. Density falls as an acute collection evolves, creating a transitional period in which the lesion may resemble the adjacent brain.

IMAGING STATE	CT RELATION TO BRAIN	INTERPRETIVE MEANING	PRINCIPAL ERROR
Acute subdural	Radiodense, appearing white	Fresh blood	Failing to distinguish an acute expanding lesion from the quieter chronic presentation
Isodense phase	Similar density to brain	Density has fallen during evolution	Calling the scan normal because collection and brain are difficult to separate
Chronic subdural	Radiolucent, appearing darker than brain	Aged blood; a contrast-enhancing membrane may border the collection	Ignoring a broad low-density extra-axial collection
Acute-on-chronic bleeding	Dense areas within a lucent region	Fresh blood within an older collection	Forcing a mixed-density lesion into a purely acute or purely chronic category

PITFALL

The isodense chronic subdural can be indistinguishable from brain tissue on density alone. A reassuring first glance must not replace a deliberate review of the extra-axial spaces, asymmetry, mass effect and any enhancing membrane, especially when the clinical story remains persuasive.

IN INDIA

Do not postpone an available CT while waiting for a theoretically preferable investigation. For the structural dementia differential, CT is generally sufficient to detect a subdural haematoma. Escalate further imaging when the initial study and the clinical picture do not agree, using the locally available radiology and neurosurgical pathway.

27.6 Why imaging is worthwhile

The yield of reversible structural disease is modest but clinically disproportionate. Structural CT or MRI detects treatable causes such as normal-pressure hydrocephalus and subdural haematoma in about **1 to 2% of demented patients**. This proportion is not the incidence of chronic subdural haematoma and must not be presented as such. It describes the group in whom imaging discovers a treatable structural explanation.

The ethical and clinical reasoning is stronger than the percentage alone. Missing a common irreversible disorder is regrettable; missing a correctable compressive lesion also forfeits recovery. Imaging therefore has a case-finding function, not merely a classificatory function. It identifies conditions in which management can change the patient's cognition, functioning, safety and future care needs.

Normal-pressure hydrocephalus, tumour, metastasis and cerebrovascular events also belong in the structural differential, but their full teaching sits in their respective sections. Their presence on the list should sharpen, not blur, the immediate question: has a subdural collection been excluded in an older person with a new headache-linked decline, subtle focal findings or relevant bleeding vulnerability?

Imaging decisions should therefore be tied to consequence as well as probability. The purpose is not to prove that every late-life cognitive complaint has a structural cause. It is to avoid declaring irreversibility before a treatable extra-axial lesion has been considered. When the history is atypical for the working psychiatric diagnosis, the decline is newly accelerated, headache is prominent, bleeding vulnerability is present or examination and history disagree, the value of CT lies in settling a question that bedside phenomenology cannot.

A report should also be reconciled with the patient. If CT demonstrates a collection, relate its side, extent, density and mass effect to the clinical syndrome with radiology and neurosurgery. If the scan is reported as unrevealing but suspicion remains high, revisit the images and the possibility of an isodense lesion rather than treating the report as infallible. Further imaging is an escalation decision, not a substitute for careful review of the initial study.

RULE **New cognitive or personality decline in an older adult demands a reversible-cause lens.** When chronic subdural haematoma is plausible, obtain structural imaging even if trauma is denied and focal neurology is slight.

27.7 Reversibility, deterioration and urgency

Correctable dementia mimic is the phrase to retain. A chronic subdural collection may resolve without an operation, and a conservative approach can be sufficient in selected cases. Other patients require surgical evacuation. Prompt diagnosis and appropriate treatment usually restore cerebral function, which distinguishes this lesion from diagnoses in which management can only slow decline or support adaptation.

Reversibility must not be mistaken for harmlessness. A small chronic extra-axial collection may exert little mass effect and cause relatively few symptoms. A large, acute or rapidly expanding

subdural haematoma can displace the brainstem and oculomotor nerve through the tentorial notch. Continued expansion may culminate in **transtentorial or foramen-magnum herniation**, an immediately life-threatening state. The full neuropsychiatric differential and management of raised intracranial pressure and hydrocephalus are covered in the dedicated section of this chapter.

Clinical deterioration changes the locus of care. Any acute decline or evidence of increasing mass effect should move the patient out of a routine psychiatric pathway and into emergency medical and neurosurgical assessment. Exact intervention thresholds cannot be improvised from phenomenology. They require current imaging, serial clinical judgement and the treating neurosurgical protocol.

> In an older patient, headache plus insidious personality or cognitive decline with few focal signs is chronic subdural haematoma until structural imaging has made the lesion implausible.

27.8 Management: LOCUS, FOCUS and MODUS

Management follows the patient's physiological risk, not the clinic in which the symptoms were first noticed. LOCUS: a stable patient with an incidentally recognised or mildly symptomatic chronic collection may enter a closely coordinated medical and neurosurgical pathway. A symptomatic, enlarging or acutely deteriorating patient requires emergency or inpatient care. Psychiatric admission alone is inappropriate when structural compression or deterioration remains active.

FOCUS: in the acute phase, protect life, recognise expansion and secure neurosurgical review. In the short term, follow headache, consciousness, behaviour, cognition, neurological findings and functional ability against the pre-illness baseline. After definitive or conservative treatment, confirm that improvement is occurring and reassess any residual psychiatric or cognitive complaint rather than assuming the original label was correct.

Recovery assessment is part of diagnosis. Improvement after the collection resolves or is evacuated supports the conclusion that it contributed materially to the decline. Incomplete recovery does not prove that treatment failed: the lesion may have coexisted with another neurological or psychiatric disorder, or the original baseline may have been inaccurately estimated. Repeat collateral history and functional assessment are therefore as important after treatment as before it. The aim is to separate recovered functions, persistent deficits and newly visible comorbidity.

MODUS: procedural treatment is surgical evacuation when the collection requires it; conservative management may suffice when the treating team judges observation appropriate, and spontaneous resolution can occur. Drug decisions focus on the medical context rather than on treating the collection with a psychotropic. Review medicines that increase bleeding tendency, but alter or reverse them only through an urgent, indication-aware medical plan. Psychological care is supportive: explain the organic basis of the symptoms, reduce frightening misattributions

and help the patient and family understand why serial observation matters. Social care addresses supervision, falls, medication administration, rehabilitation needs and safeguarding.

Care is multidisciplinary by necessity. Radiology establishes the structural evidence; neurosurgery decides whether evacuation or observation is appropriate; physicians manage comorbidity and bleeding-risk treatment; psychiatry clarifies behaviour, capacity for participation, distress and the trajectory of cognition; nursing and family observations make fluctuation or deterioration visible. At discharge, the plan should state who watches for renewed decline, who manages antithrombotic decisions and where follow-up occurs. Vague shared care is not shared care at all.

Use current institutional neurosurgical, imaging and antithrombotic guidance for operative criteria, reversal decisions and follow-up. The psychiatrist's enduring role is to recognise the mimic, avoid diagnostic closure, document change from baseline, support behaviour safely and verify cognitive recovery after the structural lesion has been treated.

28 · Brain tumours and psychiatric presentations by location

Why this matters clinically. A brain tumour may first reach psychiatry as depression, disordered thinking, personality change or cognitive impairment, even when the patient has no accompanying physical symptom. That **psychiatric-first presentation** makes structural disease easy to miss: absence of weakness, seizures or another obvious neurological complaint does not by itself establish a primary psychiatric disorder. The marks lie in recognising the possibility, interpreting location without overclaiming localisation, and separating a causal lesion from an incidental scan finding.

> **IN INDIA**
>
> Investigate a new and persistent behavioural or cognitive change in later life with structural imaging, and keep a low threshold for referral to neurosurgery or neuro-oncology.

The useful framework is anatomical rather than a memorised list of syndromes. Ask where the lesion lies, whether it compresses or infiltrates, which networks it disrupts, and whether its behavioural effects are focal or mediated by mass effect. The accompanying figure places brain tumour within the structural-lesion differential that can present to a psychiatrist.

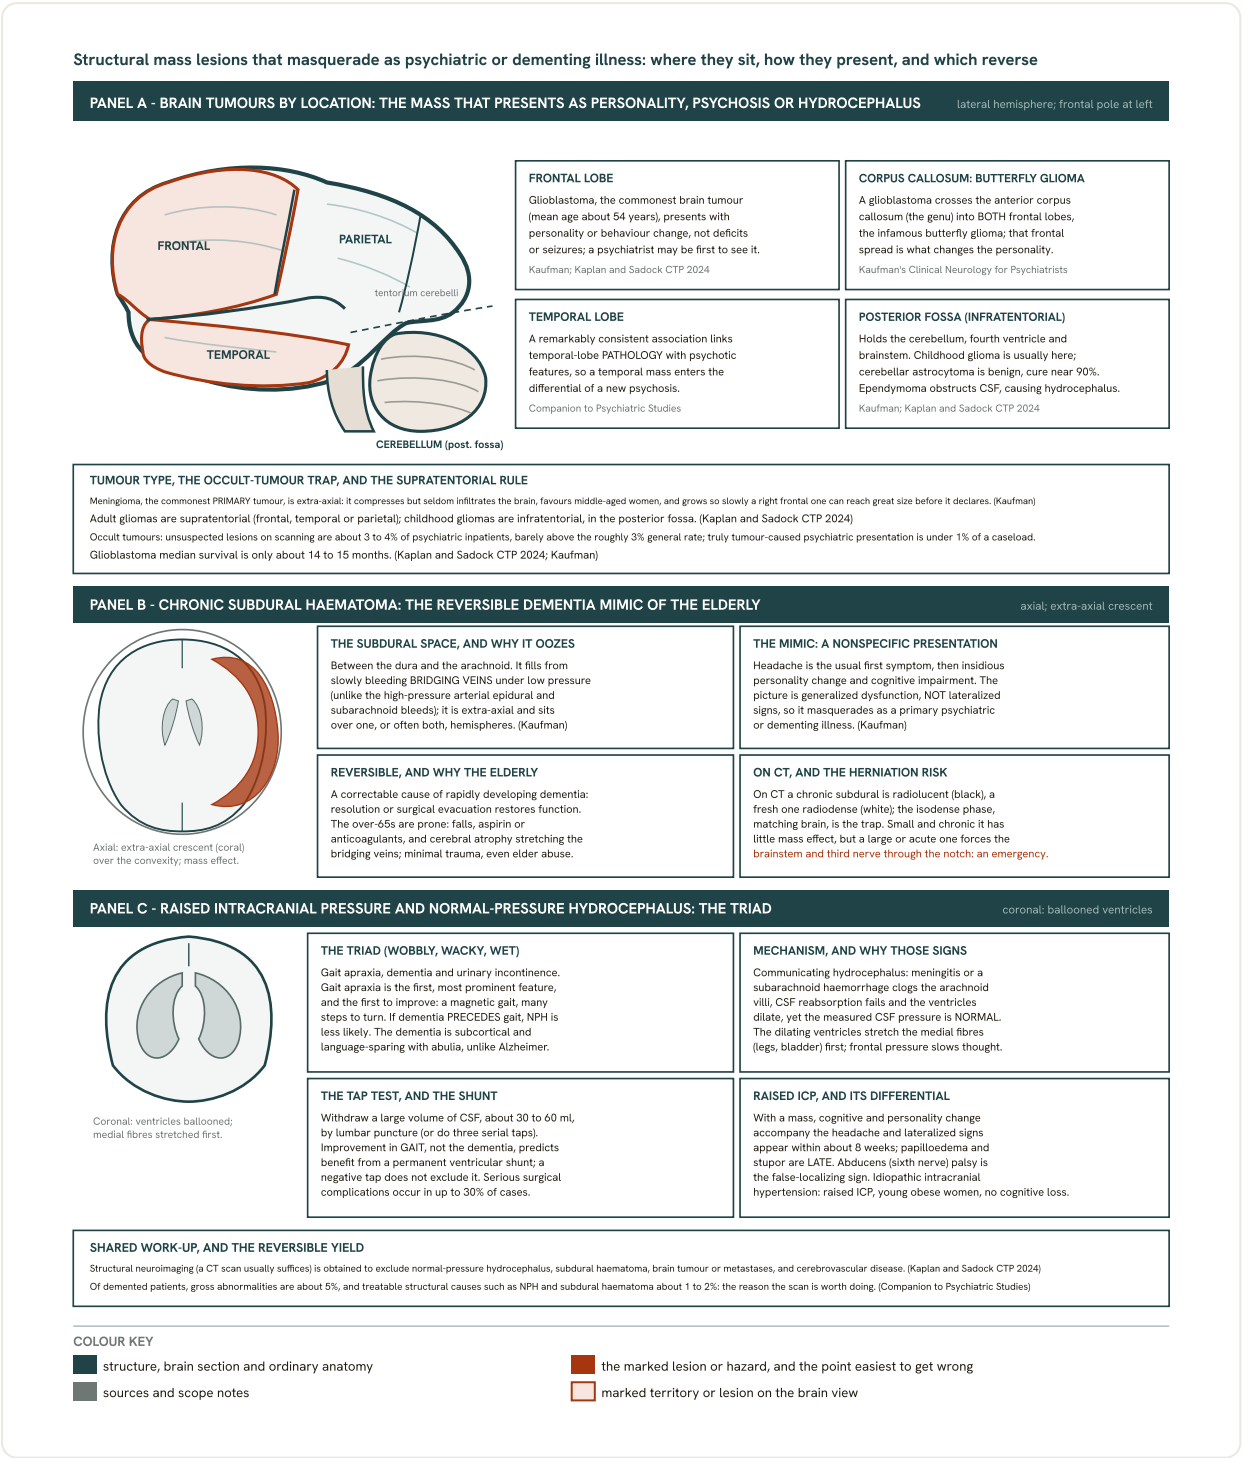


Fig 34.10 **The structural-lesion differential: masses that present as psychiatric or dementing illness.** A brain tumour is read by location: a frontal or callosal butterfly glioma changes personality, temporal-lobe pathology carries a consistent link to psychosis, and the posterior fossa is the childhood and hydrocephalus site; a chronic subdural haematoma is an extra-axial crescent from low-pressure bridging veins that masquerades, nonspecifically, as dementia in the elderly and is reversible; normal-pressure hydrocephalus is the triad of gait apraxia, dementia and urinary incontinence with ballooned ventricles at normal pressure. All three are looked for on a CT scan, which usually detects them but cannot categorically exclude an isodense subdural, a small tumour or normal-pressure hydrocephalus. (Kaufman's Clinical Neurology for Psychiatrists; Kaplan and Sadock's Comprehensive Textbook of Psychiatry 2024; Companion to Psychiatric Studies.)

RULE **A psychiatric presentation does not make the lesion psychiatrically primary.** New behavioural, affective, thought or cognitive change must be read together with tempo, examination and structural context.

28.1 The diagnostic idea: psychiatric symptoms may precede neurology

Brain tumours can mimic psychiatric illness because disturbance of cerebral tissue need not announce itself with a neatly localising deficit. Depression, thought disturbance and cognitive impairment may be the presenting problems, and the physical examination may initially appear unrevealing. The psychiatrist therefore needs a structural differential whenever the clinical story is atypical for a primary disorder. “Atypical” here is a reasoning judgment, not a single screening criterion: it emerges from the combination of change from baseline, progressive functional disruption, cognitive or personality change, and discordance between the apparent diagnosis and the course.

The causal pathways are conceptually distinct. A tumour may directly disturb tissue serving behaviour, interrupt connections between regions, compress adjacent brain, or produce more diffuse dysfunction through mass effect. These routes explain why location is informative but not deterministic. The same broad psychiatric phenotype can arise through different structural routes, while a lesion at a given site may remain clinically quiet until its effect on surrounding tissue becomes substantial.

A referral label must therefore be unpacked into phenomenology. “Depression” and “psychosis” are starting labels, not substitutes for describing affective, cognitive, behavioural, perceptual and belief phenomena. The purpose is not to relabel every atypical symptom as organic. It is to identify whether the syndrome has internal features that make a simple primary psychiatric formulation incomplete. Collateral history is particularly valuable because relatives can describe a departure from the person's established style, occupational functioning and everyday decisions when the patient experiences little change.

The review then joins phenotype to course and context. Define the change precisely, ask whether it evolves like the presumed primary disorder, seek cognitive change, headache, seizures or focal findings, and test any discovered lesion for anatomical and chronological coherence. This sequence keeps both errors in view: prematurely closing on a psychiatric diagnosis and prematurely closing on an imaging abnormality.

GOOD TO REMEMBER

When a new psychiatric syndrome is accompanied by progressive personality or cognitive change, do not allow a normal-looking brief neurological encounter to end the structural enquiry. Brain tumours may produce psychiatric disturbance without accompanying physical symptoms.

28.2 Location, compartment and biological behaviour

Location answers only part of the clinical question. In adults, most gliomas are **supratentorial**, particularly in the frontal, temporal or parietal lobes; in children, glioma is usually infratentorial in the posterior fossa. The incidence of most histological classes of primary brain tumour rises with advancing age. These observations shape the prior anatomical expectation, but they do not permit age alone to determine diagnosis.

Compartment is equally important. An intra-axial tumour arises within brain tissue and may infiltrate the networks it occupies. An **extra-axial** tumour arises outside the parenchyma and commonly acts by compression or irritation. Biological tempo then modifies the picture: a rapidly infiltrating process can disturb behaviour while physical deficits remain unobtrusive, whereas a slowly expanding compressive lesion may become large before it is recognised. The location label must therefore be interpreted alongside tissue relationship and growth pattern.

STRUCTURAL SETTING	GROUNDED ANATOMICAL POINT	PSYCHIATRIC READING
Adult glioma	Usually supratentorial, in frontal, temporal or parietal lobes	Begin above the tentorium, then localise from the full syndrome.
Childhood glioma	Usually infratentorial in the posterior fossa	Think of compartment and obstruction rather than importing an adult psychiatric localisation.
Glioblastoma	Cerebral, infiltrating, and liable to present with personality or behavioural change	Psychiatry may be the first point of medical contact.
Callosal spread	Crosses the genu into both frontal lobes	Bilateral frontal network involvement explains a behavioural presentation.
Meningioma	Extra-axial, slow-growing and compressive	Symptoms reflect the underlying brain being compressed or irritated.
Posterior-fossa tumour	May obstruct cerebrospinal-fluid flow around the fourth-ventricle compartment	Interpret psychiatric change through the hydrocephalus and raised-pressure route.

MNEMONIC

Site, Side, Structure, Symptom route

Read every tumour through its **site**, laterality or **side**, relationship to brain **structure**, and the route by which these produce the observed **symptoms**. This prevents a lobe name from becoming a diagnosis by reflex.

28.3 Frontal and callosal presentations

The strongest tumour-specific location-to-behaviour link in the available evidence is frontal involvement. Glioblastoma often presents with personality or behavioural change rather than a physical deficit or seizure, so a psychiatrist may encounter the patient before the tumour is recognised. This is not permission to reproduce the full frontal syndrome taxonomy here. Frontal lobe syndromes: dorsolateral, orbitofrontal and medial subtypes owns that localisation. The tumour question is narrower: could a structural lesion affecting frontal networks explain a new change in conduct, motivation, judgment or personality?

The classic structural pattern is the **butterfly glioma**. A glioblastoma may cross the anterior corpus callosum through the genu and infiltrate both frontal lobes. The image resembles bilateral

wings, but its clinical importance is connectional: callosal passage permits bilateral frontal involvement. Personality and behavioural change therefore follow from the frontal tissue affected, not from the visual metaphor itself.

Laterality also needs disciplined interpretation. A right frontal meningioma is a grounded example of a lesion that may become exceptionally large before causing problems. Slow expansion and the capacity of frontal regions to tolerate gradual compression can delay recognition. By the time relatives seek help, the complaint may be framed as altered character rather than neurological illness. The correct response is not to diagnose a particular frontal subtype from that history alone, but to recognise that a structural frontal process is plausible and to complete the neurological and cognitive assessment.

The interview should establish whether the change is genuinely new, and an informant account should test it against the person's previous behaviour and everyday function. A coherent cluster across personality, cognition and conduct carries more localising value than one dramatic symptom. Bilateral callosal spread should be read as a connectional and infiltrative process, not merely a midline radiological curiosity. This approach also protects against retrospective overfitting, in which every longstanding personality feature is reinterpreted after the scan.

28.4 Temporal and parietal sites: what may and may not be concluded

There is a consistent association between **temporal-lobe pathology** and psychotic features. That general neuropsychiatric association is clinically useful because it reminds the examiner that psychosis can accompany cerebral disease. It does not, however, establish that a temporal-lobe tumour characteristically produces a schizophrenia-like syndrome. The available claim is about temporal pathology in general, not tumours specifically. A candidate should preserve that scope instead of converting an association into a tumour-localisation rule.

The section on temporal-lobe syndromes provides the full focal-neuropsychiatric localisation. In this tumour fragment, temporal symptoms should prompt a structural and seizure-aware differential, but seizure classification and antiseizure-drug selection belong to the Epilepsy and psychiatry notes. A paroxysmal history may change the interpretation of unusual experiences, yet this section should not infer a seizure type or teach its pharmacology.

Parietal location must be handled with the same restraint. Adult gliomas may be situated in the parietal lobe, but the evidence does not define a distinctive psychiatric presentation for that tumour site. The section on parietal-lobe syndromes owns the focal examination. The clinically sound move is to examine for parietal dysfunction and test whether any finding aligns with the lesion, not to attach an unsupported psychiatric phenotype to the scan.

TRAP **Do not write “temporal-lobe tumours cause schizophrenia-like psychosis” from the general temporal-pathology association.** The direction from pathology to psychotic features is grounded; tumour-specific preference, symptom list and ranking are not.

28.5 Meningioma versus glioma: compression and infiltration

Meningioma arises from meningeal cells and therefore occupies an extra-axial, extraparenchymal position. Its usual effect is compression and irritation rather than infiltration of the underlying central nervous system. Meningiomas predominate in middle-aged women and often expand slowly enough to reach a large size before symptoms become evident. These features make them an excellent model for understanding why psychiatric change can be gradual, nonspecific and deceptively detached from the mass causing it.

Glioma provides the structural contrast. It is intra-axial and arises in brain tissue; glioblastoma in particular grows rapidly and infiltrates widely. The distinction matters because the scan and the syndrome must be read through different mechanisms. With meningioma, ask which cortex or network is being compressed. With an infiltrating glioma, ask which tissue and connections are being invaded, including whether spread across the corpus callosum has made the process bilateral.

A large right frontal meningioma may remain relatively silent until behavioural or functional change becomes conspicuous. This does not mean that right frontal location invariably produces one psychiatric syndrome. It means that slow extra-axial compression can conceal the biological seriousness of the lesion behind an apparently psychiatric story. Conversely, the malignant label of glioblastoma should not be inferred from behaviour alone. Personality change is a presentation clue, not a histological diagnosis.

The comparison also explains why size alone is an unreliable proxy for symptom severity. A slowly compressive lesion can permit functional adaptation until reserve is exceeded, whereas an infiltrating lesion may disrupt critical networks without producing an equally impressive external contour. Clinical interpretation therefore requires the relationship between lesion and brain, not simply a measurement of the mass. For psychiatry, the practical question is whether the structural process offers a better account of the new syndrome than a coincidental scan finding does.

> **PITFALL**
>
> Do not treat every lesion on imaging as causal merely because the patient has psychiatric symptoms. Match the timing, anatomical effect and clinical course. A slow extra-axial mass, an infiltrating intra-axial tumour and an incidental abnormality carry different causal arguments.

28.6 Posterior fossa and the obstructive route

The **posterior fossa** contains the cerebellum, pons, medulla, fourth ventricle, and vertebral and basilar arteries. Childhood gliomas are usually infratentorial, and childhood cerebellar astrocytomas tend to be benign, cystic and noninvasive. Complete removal can be achieved with a cure rate approaching **90%**. This paediatric anatomical pattern should not be extrapolated into a distinct adult psychiatric syndrome without evidence.

The grounded psychiatric route is obstruction. A tumour arising near the ventricular cerebrospinal-fluid pathway may impede flow and produce hydrocephalus; ependymoma is a

supported example. Mental or behavioural change in this setting is therefore interpreted as a possible consequence of hydrocephalus or raised intracranial pressure, not as proof of an independent cerebellar cognitive-affective syndrome. Raised intracranial pressure, hydrocephalus, and the psychiatric implications of post-injury seizures is taught separately and owns the clinical syndrome, differential and seizure-overlap details.

This boundary is important in examination answers. “Posterior fossa” is an anatomical compartment, not a psychiatric diagnosis. A strong answer identifies its contents, recognises the childhood infratentorial distribution, explains the obstruction route, and stops before supplying an unsupported adult symptom complex. In practice, the same discipline prevents overinterpretation of emotional or cognitive symptoms when the decisive mechanism may be diffuse pressure-related cerebral dysfunction.

28.7 Occult lesions, incidental findings and causal attribution

Structural disease deserves vigilance, but its base rate must be interpreted accurately. The accompanying strip shows that clinically unsuspected lesions in psychiatric inpatients are only slightly more frequent than incidental scan abnormalities in the general nonpsychiatric population. Patients whose psychiatric symptoms are actually caused by a hidden brain tumour remain a vanishing minority across a general psychiatrist's career. The comparison warns against assuming that every unexpected lesion explains the psychiatric presentation.

3%-4%	~3%	<1%
UNSUSPECTED LESIONS IN PSYCHIATRIC INPATIENTS	INCIDENTAL SCAN ABNORMALITIES IN THE GENERAL POPULATION	TRUE HIDDEN-TUMOUR PRESENTATIONS IN A PSYCHIATRIST'S CAREER CASELOAD

The diagnostic problem is therefore asymmetric. Missing a causal structural lesion can have grave consequences, but labelling an incidental lesion as causal can also distort care. The solution is neither indiscriminate scanning nor reassurance based only on rarity. It is structured attribution: determine whether the clinical course is compatible with the lesion, whether the anatomical site plausibly relates to the observed change, whether the imaging effect is substantial, and whether competing psychiatric or medical explanations remain stronger.

Base rates are especially useful after, not instead of, clinical assessment. Before imaging, they prevent the rare hidden tumour from becoming the default explanation for every late or unusual psychiatric presentation. After imaging, they prevent the mere presence of an abnormality from being treated as proof of cause. The posterior probability becomes persuasive only when the lesion and the syndrome converge anatomically, temporally and functionally. This is why an incidental lesion in a well-characterised primary psychiatric illness and an infiltrative frontal mass accompanying progressive personality change demand different interpretations even though both scans are “abnormal”.

- **Concordance:** does the behavioural or cognitive change fit the tissue affected without requiring an extravagant localisation claim?

- **Chronology:** did the syndrome and functional decline evolve in a way compatible with the lesion?
- **Burden:** is there compression, infiltration, bilateral spread or obstruction that could plausibly disturb brain function?
- **Alternatives:** could the scan finding be incidental while another disorder better explains the presentation?

The practical balance. Maintain a low threshold for reconsidering structure when the course is progressive or the formulation is unstable, but do not turn imaging into a substitute for phenomenology, collateral history and examination. The bedside cognitive examination as a general method is taught elsewhere in this chapter, while standardised instruments and their scoring belong to the Psychotherapies, clinical assessment, MSE and nosology notes.

28.8 Glioblastoma as the high-stakes psychiatric presentation

Glioblastoma is described in the cited neurology source as the most malignant astrocytoma. It develops almost exclusively in the cerebrum, grows rapidly and infiltrates widely. Crucially for psychiatry, personality or behavioural change may precede physical deficits or seizures, placing the psychiatrist at the front of the diagnostic pathway.

Kaufman's Clinical Neurology for Psychiatrists gives a mean age at diagnosis of **54 years**. That source-specific value should be recognised as such rather than silently replaced by a remembered epidemiological figure. Prognosis is poor: a field-leading trial reported median survival around **14 months**, with about one-quarter alive at 2 years and only 10% alive at 5 years. Kaufman's rounded textbook estimate is about 15 months; the close figures should remain separately attributed rather than falsely harmonised.

The examination answer should finish by linking biology to presentation. Cerebral infiltration can accompany psychiatric-first illness; frontal involvement explains prominent personality and behaviour change; and callosal spread may create the bilateral frontal butterfly pattern. None of these features permits psychiatry alone to determine histology. They identify when a structural explanation deserves urgent investigation and multidisciplinary neurological care.

A concise synthesis should retain uncertainty where the source does. Brain tumour belongs in the differential of an unexplained psychiatric change, but most psychiatric patients with an unexpected structural finding do not thereby become hidden-tumour cases. Conversely, rarity must not neutralise a coherent progressive syndrome. The mature answer holds both propositions together: localisation guides suspicion, and causal attribution requires concordance.

> Memorise the logic, not an unsupported localisation table: psychiatric-first presentation is possible, frontal and callosal glioblastoma has a grounded behavioural link, the general temporal-pathology association does not establish temporal-tumour psychosis, and posterior-fossa psychiatric change is taught through obstruction and pressure.

29 · Raised intracranial pressure, hydrocephalus, and the psychiatric implications of post-injury seizures

The psychiatric presentation may precede the obvious neurological emergency. Raised intracranial pressure, hydrocephalus and post-injury seizures belong together because each can first appear as altered behaviour, slowed cognition, confusion or apparent psychiatric relapse. The scoring points lie in recognising the structural pattern, distinguishing pressure syndromes from normal-pressure hydrocephalus, and interpreting an episodic post-injury mental-state change as potentially epileptic rather than automatically psychogenic or psychiatric.

> **IN INDIA**
>
> When a mass or raised pressure is suspected, obtain structural imaging before attributing a new presentation to a primary psychiatric cause; phenytoin and levetiracetam are reasonable choices for post-injury seizures.

This section therefore follows a psychiatrist's route through the problem: identify pressure-related danger, recognise the treatable hydrocephalus phenotype, classify post-traumatic events by their relationship to injury, and prescribe psychotropics without casually worsening seizure vulnerability. The full classification of seizures and antiseizure-drug pharmacology belong to *the Epilepsy and psychiatry notes*. The dementia syndrome and its general diagnostic work-up belong to *the Dementias and neurocognitive disorders notes*.

Late PAPILLEDEMA AND STUPOR	**30-60 mL** HIGH-VOLUME TAP	**>7 days** PTE WINDOW	**Up to 50%** MAJOR-TBI PTE BURDEN

29.1 Raised intracranial pressure: the psychiatric warning pattern

Raised intracranial pressure from a mass lesion need not announce itself with dramatic neurological signs. Headache may be accompanied by subtle cognitive and personality change, while lateralised findings usually emerge within **8 weeks**. The clinically dangerous misconception is that a normal-looking fundus or preserved alertness early in the course removes a mass lesion from consideration. Papilloedema and stupor are overt signs that generally appear late, if they appear at all.

Once enough blood or mass has accumulated, headache, confusion, nausea, vomiting and focal neurological signs may accompany the pressure rise. For the psychiatrist, the sequence matters more than any isolated complaint. A new headache linked to progressive cognitive or personality change, followed by vomiting or focal findings, is a structural story until investigated. The absence of a fully developed pressure syndrome should not be mistaken for evidence against intracranial disease.

Structural imaging is used to exclude normal-pressure hydrocephalus, subdural haematoma, tumour or metastasis, and cerebrovascular lesions as contributors to cognitive or behavioural change; CT is generally sufficient for this screening purpose. This is especially important when the history has been framed as depression, dementia, medication non-adherence or personality

change. Imaging does not replace neurological examination, but it prevents psychiatric phenomenology from being asked to settle an anatomical question it cannot settle.

GOOD TO REMEMBER

Cognitive and personality change can accompany headache before papilloedema or stupor appears. A late overt sign is not a sensible prerequisite for early imaging and referral.

29.2 False localisation and the raised-pressure differential

Pressure can distort structures distant from the causative lesion. A unilateral or bilateral **abducens nerve palsy** may result when raised pressure stretches and damages the sixth cranial nerve. The affected eye turns inward because medial rectus action supplied by the intact third cranial nerve is unopposed. This is a false-localising sign: it signals global pressure mechanics and does not by itself prove a pontine lesion at the side of the ocular deviation.

PITFALL

Do not localise every sixth-nerve palsy directly to the pons. Raised intracranial pressure can stretch either or both abducens nerves, so the ocular sign must be interpreted with headache, cognition, vomiting, fundus findings and imaging.

Idiopathic intracranial hypertension provides a crucial contrast because pressure is raised without the cognitive and seizure pattern expected from many structural lesions. It occurs predominantly in young obese women with menstrual irregularity; elevated venous-sinus pressure, often related to sinus occlusion or stenosis, is implicated in many cases. The headache is dull and generalised and may accompany pulsatile tinnitus or transient visual obscurations. Papilloedema is the distinctive finding and can progress through an enlarged blind spot to blindness from optic atrophy if untreated.

CLINICAL SETTING	PSYCHIATRIC OR COGNITIVE CLUE	NEUROLOGICAL DISCRIMINATOR	CLINICAL IMPLICATION
Mass-related raised pressure	Subtle cognitive or personality change may accompany headache	Papilloedema and stupor are usually late or absent; focal signs may emerge	Do not wait for the overt syndrome before imaging.
Idiopathic intracranial hypertension	Cognitive capacity is not affected and seizures are not produced	Papilloedema with visual obscurations or pulsatile tinnitus	Protect vision and investigate the pressure disorder rather than attributing symptoms to a mass.
Normal-pressure hydrocephalus	Cognitive decline occurs with gait apraxia and urinary incontinence	Ventricular enlargement occurs despite normal measured CSF pressure	Recognise a potentially correctable cause of cognitive decline.

29.3 Normal-pressure hydrocephalus: syndrome and mechanism

Normal-pressure hydrocephalus is defined clinically by dementia, urinary incontinence and gait apraxia and is important because it is a correctable cause of cognitive decline. The triad should be stated by content. An unsupported eponym adds no diagnostic value and is best omitted.

The cognitive pattern is **subcortical**: thought and gait slow, language is relatively spared, and abulia is common. The patient may therefore appear apathetic, poorly motivated or depressed, but the accompanying motor and bladder phenotype changes the diagnostic frame. The physical features help separate this syndrome from cortical dementias, although formal dementia criteria and the complete general work-up remain outside the present section.

Most cases are idiopathic, while meningitis or subarachnoid haemorrhage may precede secondary cases. In the latter setting, inflammatory material or blood can obstruct cerebrospinal-fluid reabsorption at the arachnoid villi. Production continues, fluid accumulates and the ventricles expand. Measured CSF pressure and CSF protein and glucose concentrations remain normal. The label therefore describes the pressure measured during evaluation, not the absence of a pathophysiological disturbance in CSF circulation.

Anatomy explains the clinical triad. Ventricular expansion compresses brain parenchyma and stretches corticospinal and other internal-capsule tracts. The most medial fibres supply the legs and voluntary bladder musculature, making gait and urinary symptoms prominent. Pressure on frontal lobes contributes to cognitive impairment and psychomotor slowing. The syndrome is thus a unified periventricular and frontal network disorder, not a collection of unrelated complaints.

MNEMONIC

Walk, think, void

Walk: gait apraxia is usually earliest and most prominent. **Think:** cognition slows with language relatively spared. **Void:** urgency and frequency may progress to incontinence. The order keeps gait at the front of diagnostic reasoning.

29.4 Gait-first recognition, tap testing and shunting

The **magnetic gait** is the hallmark. Gait apraxia is generally the initial, most consistent and most prominent feature, and it is the first feature to improve with treatment. Patients may require multiple small steps to turn, yet can often step over a stick because the stepping reflex is relatively preserved. If dementia clearly precedes gait dysfunction, normal-pressure hydrocephalus becomes less likely. Bladder symptoms often evolve from urgency and frequency to incontinence and may also improve after treatment.

The high-volume tap test asks a practical question: does temporary CSF removal improve the component most likely to respond? Withdrawal of **30 to 60 mL** by lumbar puncture, or serial lumbar punctures, is followed by reassessment of gait. Improvement in gait, not necessarily cognition, supports the diagnosis and predicts benefit from permanent drainage. A negative test

does not exclude the diagnosis. This asymmetry is worth stating clearly: a positive gait response is informative, whereas absence of response is not a definitive rule-out.

Permanent drainage uses a shunt from a lateral ventricle to the chest or abdominal cavity. Benefit remains difficult to predict, and serious surgical complications occur in up to 30% of cases. The decision therefore requires neurological and neurosurgical assessment rather than an automatic transition from triad to operation. The psychiatrist contributes reliable baseline and post-tap descriptions of gait, cognition, motivation and function, while the procedural team weighs likely gain against harm.

RULE **In suspected normal-pressure hydrocephalus, follow the gait.** It is usually the first feature to appear and the first to improve; gait improvement after CSF removal predicts shunt benefit more usefully than immediate cognitive change.

29.5 Post-traumatic seizures: substrate, timing and risk

Post-traumatic epilepsy arises because a cortical scar at the injury site, and sometimes retained foreign material, forms an epileptogenic focus. The onset is focal and localises ipsilaterally to the lesion, with possible secondary generalisation; after traumatic brain injury the usual description is complex partial seizures with secondary generalisation. This anatomical logic matters in psychiatry because a stereotyped behavioural spell may be the expression of a focal cortical discharge rather than a new primary psychiatric syndrome.

Timing separates provoked post-traumatic seizures from epilepsy. **Immediate and early post-traumatic seizures** occur within the first 24 hours and within the first 7 days respectively and are classified as provoked, not as post-traumatic epilepsy. Only recurrent seizures **more than 7 days** after injury meet the stated temporal definition of post-traumatic epilepsy. These trauma-specific boundaries must not be copied into stroke without a separate source.

Burden follows injury biology. After major traumatic brain injury, prevalence rises with time and may eventually reach **50%**; minor head trauma rarely produces post-traumatic epilepsy. Risk is greater after penetrating than after blunt or blast injury, greater with alcohol abuse, and increases with traumatic-brain-injury severity. It may account for 5% to 6% of all epilepsy, and remission occurs in only 25% to 50%. These figures describe different denominators and should never be collapsed into one prevalence statement.

- **Substrate:** a cortical scar at the lesion site generates focal-onset seizures.
- **Timing:** events in the immediate and early windows are provoked; recurrent events after the early window define post-traumatic epilepsy.
- **Risk:** penetration, alcohol abuse and greater injury severity raise the likelihood.

29.6 Psychiatric implications after traumatic brain injury

Post-traumatic epilepsy adds more than recurrent convulsions to an already injured brain. It causes disability and creates a risk of further head injury, while antiseizure drugs may worsen trauma-related cognitive and personality change. Treatment review must therefore distinguish injury effects, seizure effects, medicine effects and a comorbid psychiatric disorder. A new

complaint of slowing, irritability or altered behaviour after treatment begins should not be assigned to "the TBI" without examining this layered causation.

Psychogenic nonepileptic seizures also enter the differential because some studies report increased PNES among people with milder traumatic brain injury. This does not justify diagnosing PNES merely because the injury was mild or investigations are unrevealing. It means that recurrent post-injury events need careful semiology and appropriate neurological assessment, while psychological formulation remains available when supported. The distinction is consequential: mislabelling epilepsy as psychogenic withholds seizure care, while mislabelling PNES as epilepsy exposes the patient to ineffective medication and its cognitive burden.

Seizures also interact with the broader psychiatric aftermath. Post-traumatic depression occurs in 25% to 50% of patients with traumatic brain injury, correlates only weakly with injury severity, and is associated with poor quality of life; post-traumatic epilepsy is one of the trauma-related factors that further increases its incidence. The depression syndrome itself belongs in the dedicated section on psychiatric sequelae of traumatic brain injury. Here the point is causal layering: seizure burden can amplify disability, occupational loss and social restriction, all of which alter presentation and recovery.

GOOD TO REMEMBER

A one-week course of phenytoin or levetiracetam after severe traumatic brain injury reduces early seizures only during that period; no antiseizure drug clearly prevents later post-traumatic epilepsy. Prophylaxis for the early window is not disease modification.

29.7 Post-stroke seizures and behavioural presentations

The epileptogenic substrate after stroke is likewise a cortical scar, producing focal-onset seizures that may secondarily generalise. In adults older than 65 years, approximately 50% of newly developing seizures originate from stroke scars. This is the proportion of later-life new seizures attributable to stroke scars, not the proportion of all stroke survivors who develop seizures.

A first seizure after 60 years may herald a cerebral tumour, but stroke is approximately equally likely; in either setting seizures typically begin focally and then generalise. Structural imaging and neurological assessment therefore matter even when the event has resolved and the mental state has normalised. Tumour presentations by location are covered elsewhere in this chapter, and seizure classification remains with *the Epilepsy and psychiatry notes*.

A cortical-stroke survivor who develops confusion, unresponsiveness, abnormal movements or unusual behaviour requires seizure in the differential. This is the safety bridge to psychiatry: episodic or unexplained behavioural change can be ictal or postictal, including when a major convulsion is not witnessed. The working formulation should stay neurological until the temporal pattern, examination and appropriate investigations have been reconciled.

Seizures may also participate in secondary psychosis after stroke. In one small lesion-matched comparison, seizure disorder was present in three of five patients with post-stroke psychosis and none of five stroke controls without psychosis. Together with seizures and subcortical atrophy,

right-hemisphere lesions, especially right lateral prefrontal involvement or disrupted connectivity to that area, appear important in pathogenesis. The tiny comparison supports an association, not a diagnostic rule or causal certainty.

29.8 Seizure-aware psychiatric prescribing: locus, focus and modus

Locus. A patient with an acute seizure, impaired recovery, new focal sign or pressure-related deterioration belongs in emergency or inpatient neurological care, not routine psychiatric follow-up. Stable longer-term psychiatric care can proceed collaboratively as an outpatient when the seizure disorder and structural lesion are being monitored. **Focus.** In the acute phase, preserve neurological safety and avoid obscuring evolving mental-state change. Over subsequent care, target disabling mood, psychotic or behavioural symptoms while repeatedly reviewing seizure control, cognition, personality and adverse effects. **Modus.** Combine appropriate neurological treatment, cautious psychotropic use, psychological work, injury prevention and social rehabilitation rather than treating medication selection as the whole plan.

Antipsychotics require caution in post-stroke psychosis because stroke increases seizure vulnerability and antipsychotics can lower the seizure threshold. Risk is agent-specific and is magnified by rapid introduction and high doses. Haloperidol, risperidone, paliperidone and aripiprazole are among the least threshold-lowering agents, whereas clozapine, olanzapine, quetiapine, chlorpromazine, thioridazine and perphenazine are among the more concerning agents. Clozapine-associated seizures occur in **1% to 4.4%** of treated patients, particularly with rapid dose escalation.

TRAP **"Atypical antipsychotics are safer for seizures" is not a reliable rule.** The epilepsy-focused table places clozapine, olanzapine and quetiapine in the high-potential group alongside chlorpromazine, thioridazine and perphenazine, while haloperidol, risperidone and aripiprazole are low-potential choices. Another chapter offers a broader class-level reassurance, so retain the conflict and answer with agent, dose and titration rate rather than generation.

PRESCRIBING QUESTION	LOWER-CONCERN POSITION	HIGHER-CONCERN POSITION	PRACTICAL READING
Antipsychotic choice	Haloperidol, risperidone, paliperidone, aripiprazole	Clozapine, olanzapine, quetiapine, chlorpromazine, thioridazine, perphenazine	Avoid rapid introduction and high doses.
Antidepressants at therapeutic use	Therapeutic doses usually do not promote seizures; SSRIs and SNRIs are tabulated as low potential	Overdose or intoxication with tricyclic, tetracyclic, SSRI or SNRI agents can cause seizures	Separate ordinary treatment from overdose risk.
Particular antidepressants	Choose according to the whole psychiatric and neurological formulation.	Amoxapine, bupropion, clomipramine and maprotiline are usually avoided in epilepsy	Do not generalise from antidepressant class alone.
Drug interactions	Some newer antiseizure drugs have relatively little enzyme induction or inhibition	Antiseizure drugs may reduce psychotropic efficacy; their withdrawal may raise psychotropic levels; psychotropics may raise antiseizure-drug levels to toxicity	Review each direction whenever either regimen changes.

When a psychotropic is necessary, start at a low dose, titrate slowly, and monitor antiseizure-drug levels and EEGs where clinically indicated. Current neurological guidance should be consulted for the underlying seizure disorder because the allowed evidence does not supply a named guideline. Full antidepressant selection belongs to *the Mood disorders: pharmacotherapy notes*, while detailed antiseizure-drug interactions and drug choice belong to *the Epilepsy and psychiatry notes*.

- **Before prescribing:** clarify whether recent events are controlled epilepsy, provoked seizures or an unresolved episodic-behavioural syndrome.
- **During titration:** avoid rapid escalation and unnecessarily high exposure, and track cognition and personality as well as symptom response.
- **After any regimen change:** reconsider bidirectional metabolic interactions and the possibility of toxic or ineffective levels.

The must-memorise line: gait usually comes first and improves first in normal-pressure hydrocephalus; early post-traumatic seizures are provoked, whereas recurrent seizures after 7 days define post-traumatic epilepsy; and psychotropic seizure risk is agent-, dose- and titration-dependent.

Consolidate and apply

30 · Worked vignettes

Vignettes reward controlled inference, not instant recognition. The examiner gives a compressed story in which psychiatric symptoms, cognitive change, neurological findings and medical context compete for attention. Marks come from showing how the clues alter urgency, syndrome formulation, localisation, differential diagnosis and disposition. A disease label without that chain looks guessed, even when it happens to be correct.

These cases are deliberately qualitative. Their purpose is to demonstrate how knowledge from the disease-specific sections is used without repeating diagnostic thresholds, lesion tables, instrument scores or drug regimens. In every vignette, preserve uncertainty long enough to identify the dangerous alternative, the reversible contributor and the observation that would most efficiently change the formulation.

30.1 Read the stem as a changing clinical state

Vignette compression means that every included clue has a job, but not every clue has equal weight. Begin by converting the narrative into a brief problem representation: previous function, nature of change, tempo, setting, dominant psychiatric or cognitive feature, relevant neurological context and present safety concern. Do not repeat the stem. Reorganise it into information that supports a decision.

A useful reading separates observations from interpretations. “Family reports that the patient no longer manages familiar routines” is an observation about changed function. “The patient has dementia” is an interpretation requiring a broader syndrome assessment. “The patient laughs during serious conversation” is an observation. Depression, disinhibition, emotional incontinence and socially incongruent coping are competing interpretations. This distinction protects against making the opening adjective in the stem carry the whole diagnosis.

State	Tempo	Pattern	Risk
ESTABLISH	ORDER	TEST	ACT

Next, mark clues that change action rather than merely increase diagnostic confidence. A deteriorating level of engagement, new focal finding, recent injury, seizure-like event or pressure-related concern may require urgent medical escalation before refined psychiatric classification. Then identify the negative information that is conspicuously absent. Missing collateral history, an untested level of arousal or an unexplored medication change often represents the next task rather than evidence that nothing is wrong.

Read the command in the question as carefully as the stem. “Most likely diagnosis” requires a concise choice with discriminating support. “Discuss assessment” requires a sequence that moves from safety and state to focused examination and investigation. “Outline management” requires setting, targets, modes of care and monitoring. Do not answer the disease you hoped to see. Answer the decision the vignette actually asks you to make.

RULE **Translate every important clue into a clinical consequence.** State whether it changes urgency, syndrome, localisation, cause, investigation or management. A clue that changes none of these should not dominate the answer.

30.2 The quiet inpatient who is "just depressed"

A patient receiving medical care becomes withdrawn, eats poorly, speaks little and seems intermittently disengaged. Staff describe sadness, while relatives insist that the behaviour is new and variable. The opening task is not to choose between depression and delirium from surface appearance. It is to establish **state**: arousal, ability to sustain engagement, reliability of the interview and departure from the patient's usual functioning.

The answer should explicitly resist equating low activity with a primary mood syndrome. Reconstruct the course from nursing observations and family collateral, review the medical trajectory and treatment exposures, and perform a bedside assessment that the patient can actually understand and complete. If engagement varies, serial observations are more informative than a polished snapshot. A quiet patient may create less disruption while carrying greater medical risk than an agitated patient.

The case may ultimately fit delirium, depression, medication effect, physiological compromise or a mixed picture. The dedicated sections on delirium provide its criteria, screening operations, causes, prevention and treatment. The vignette skill is to show why the syndrome must be considered and what must happen before a psychiatric conclusion is safe. If substance withdrawal becomes plausible, name it and cross-refer to the Substance use disorders notes rather than reproducing its timeline or regimen.

PITFALL

Do not let ward labels substitute for examination. "Confused," "uncooperative," "depressed" and "sleepy" are summaries made by observers. Convert each label back into behaviour, trajectory and testable mental-state findings.

A complete response also states disposition. Physiological instability, worsening cognition, unreliable supervision or inability to participate in essential care supports a medical locus rather than routine psychiatric follow-up. Symptom control is subordinate to identifying and treating the underlying process.

State what would make you revise the formulation after initial care. Improvement with correction of the medical disturbance may support the working account, but persistence should trigger reassessment rather than retrospective certainty. A mixed presentation remains possible, so continuing mood symptoms can be evaluated more reliably after attention, arousal and physiological status have stabilised.

30.3 Emotional change after a vascular event

A patient recovering from a vascular event is described as tearful, disengaged and difficult to motivate. Family members ask whether this is depression, while staff interpret the behaviour as

poor cooperation. This is a **syndrome competition** vignette. The candidate should not collapse tearfulness, reduced initiation, fatigue, subjective distress and loss of function into a common label.

Clarify what the episodes look like, what triggers them, whether the expressed emotion matches the sustained internal mood, and whether reduced activity reflects desire, initiation, physical limitation, cognitive burden or fear. Ask how the patient understands the disability and whether awareness of deficit is intact. Examine attention and communication before interpreting apparent indifference. Language impairment can hide distress, while executive dysfunction can impair initiation despite preserved concern.

The differential may name post-stroke depression, apathy, fatigue, emotional incontinence, anxiety, psychosis and denial of deficit, but each belongs to its disease-specific section. In the vignette answer, use them as hypotheses and state the discriminating question or observation rather than teaching their epidemiology, lesion correlates or treatment evidence. Antidepressant pharmacology and comparative drug selection belong to the Mood disorders: pharmacotherapy notes.

The formulation should remain layered. Direct effects of brain injury, psychological response to disability, pain, sleep disruption, medication effects and pre-existing psychiatric vulnerability may coexist. Explain which account currently leads, what competing account remains clinically important, and what collateral or examination finding would change the ranking. Avoid claiming that emotional phenomenology uniquely identifies lesion location.

In an examination answer, pair each hypothesis with its next discriminator. Sustained mood change calls for a full mood assessment; episodic emotional expression calls for careful phenomenology; reduced goal-directed behaviour calls for assessment of initiation, cognition and environmental demand. This format demonstrates comparison without stealing the detailed criteria or treatment material from the owning sections.

GOOD TO REMEMBER

When relatives say “personality has changed,” ask for examples from comparable situations before and after the event. Concrete changes in initiation, inhibition, empathy, judgement or emotional expression are more useful than a global character judgement.

30.4 Behavioural disturbance after head injury

After a head injury, a patient becomes irritable, impulsive and forgetful. The family requests medication because the behaviour is disrupting care. The tempting answer is a list of drugs. The stronger answer begins with **baseline shift**: document the pre-injury pattern, the injury course, current cognitive state, environmental triggers, sleep and pain, substance exposure, family interaction and the functional consequences of the change.

Separate the original injury from active complications and later psychiatric sequelae. Current agitation may reflect confusion, fear, communication difficulty, overstimulation, frustration, disinhibition, a mood or psychotic syndrome, substance-related factors, or a new medical

problem. Memory complaints require clarification of attention, acquisition, retrieval, awareness and everyday function. Formal injury classification and post-traumatic amnesia belong to their dedicated sections; use those constructs in the formulation without restating thresholds.

The immediate plan follows locus, target and mode. Decide whether behaviour can be managed safely in the present setting or requires higher medical supervision. Define the target behaviour precisely, identify antecedents and consequences, modify the environment, support communication, involve rehabilitation and caregivers, and consider medication only after the reason for using it is clear. Pharmacological management in brain injury has its own section because vulnerability to adverse cognitive, behavioural and neurological effects changes prescribing judgement.

A good answer includes measurable clinical outcomes without inventing a scale: fewer dangerous episodes, improved participation in care, more stable sleep, safer mobility, better communication and reduced caregiver strain. Reassess whether apparent improvement reflects genuine recovery or merely sedation. If seizure classification or anticonvulsant choice becomes relevant, cross-refer to the Epilepsy and psychiatry notes.

Causal language should remain proportionate. Temporal association with injury supports relevance but does not prove that every later symptom is a direct lesion effect. A strong formulation distinguishes direct neurological consequences, treatment and environmental contributors, psychological responses to disability, and independent psychiatric comorbidity. It then assigns each contributor a corresponding intervention rather than seeking a universal explanation.

30.5 The confident but inaccurate history

A patient gives fluent, confident accounts that conflict with collateral history, appears unaware of the discrepancies and cannot reliably retain recent information. The error is to call the patient psychotic or deceptive from the inaccuracy alone. Begin with **memory monitoring**: determine whether the problem lies in attention, learning, retention, retrieval, temporal placement, source attribution or checking the accuracy of what is recalled.

Observe how the account changes with open questions, focused prompts and correction. Ask whether the patient is filling gaps, repeating familiar narratives, misplacing genuine events or defending a fixed belief despite contrary evidence. Evaluate language comprehension, executive control and awareness, because each can shape the narrative. Collateral history is not merely corroboration here. It is part of the examination of memory and function.

The differential can name an amnestic syndrome, confabulation, delirium, a focal neurological process, a seizure-related episode and a broader neurocognitive disorder. The dedicated amnestic sections provide the pathology and distinctions. Emergency nutritional disease and repletion belong to the Endocrine and metabolic psychiatry notes. The full dementia syndrome and its diagnostic workup belong to the Dementias and neurocognitive disorders notes.

Management must account for the practical consequences of unreliable memory: consent discussions, treatment adherence, wandering, vulnerability, financial decisions and caregiver burden. Do not confront inaccuracies as moral failures. Provide orientation and external memory support, simplify communication, treat the underlying cause and reassess capacity for the decision at hand. Standardised cognitive instruments may assist assessment and follow-up, but their content and scoring belong to the Psychotherapies, clinical assessment, MSE and nosology notes.

Capacity should be framed around the actual decision, not inferred globally from poor recall. Present information in an accessible form, check understanding and retention relevant to that decision, and seek support where impairment prevents safe independent action. The same patient may participate meaningfully in some choices while requiring protection or assistance for more complex consequences.

30.6 A focal error hidden inside a psychiatric complaint

A patient is referred for odd behaviour after repeatedly missing objects, mishandling familiar actions or becoming unable to complete previously routine tasks. Conversation remains socially fluent, encouraging a purely psychiatric explanation. The correct pivot is **convergent localisation**: ask whether several observations point toward a coherent cognitive network while controlling for vision, hearing, language, motor ability, attention and comprehension.

Describe the failed operation before naming a lobe syndrome. Determine whether the patient detected the stimulus, understood the command, recognised the object, planned the action and could execute the required movement. Compare spontaneous behaviour with formal testing and use another sensory or response modality when possible. A task failure is not anatomically pure; it reflects every process required between instruction and response.

VIGNETTE PATTERN	REASONING PIVOT	UNSAFE SHORTCUT
Fluent conversation with failed daily tasks	Test specific cognitive operations and function	Assuming social fluency means intact cognition
Apparent inattention to part of space	Separate spatial attention from sensory loss and unawareness	Calling every omission poor effort
Inability to perform a familiar action	Control for comprehension, weakness and coordination	Naming a praxis disorder from a contaminated task
Failure to identify a familiar stimulus	Confirm adequate sensation and language access	Equating recognition failure with memory loss
Disinhibited or apathetic behaviour	Seek executive, social and motivational convergence	Localising from personality description alone
Episodic unusual behaviour	Reconstruct onset, offset, awareness and recovery	Assuming psychiatric causation because examination later looks normal

TRAP **A memorable sign is not a complete localisation.** The examination answer must show convergence, exclude task contamination and acknowledge that distributed networks can produce mixed findings.

After establishing the pattern, name the relevant frontal, temporal, parietal, visual-association or disconnection syndrome and direct the reader to its dedicated section. The vignette marks lie in demonstrating why the label fits and what competing explanation has been tested, not in reproducing the owner section's feature list.

Keep psychiatric phenomenology in the formulation even when localisation appears convincing. Distress, suspiciousness, demoralisation, impulsivity or altered emotional expression may arise through several pathways and can shape risk independently of the focal deficit. Conversely, a psychiatric history does not invalidate a reproducible neurological pattern. The task is integration, not choosing which specialty owns the patient.

30.7 Slowly evolving psychiatric change with neurological risk

A patient develops progressive slowing, irritability, reduced initiative or impaired judgement. The referral asks for treatment of a primary psychiatric disorder, but collateral history suggests a clear departure from baseline and emerging neurological or functional change. This requires **structural vigilance**: preserve structural, pressure-related, vascular and other neurological explanations until tempo, examination and appropriate investigation make them less likely.

Look for a trajectory that does not fit the opening formulation. Ask about headache pattern, vomiting, gait, falls, injury, seizure-like events, focal symptoms, continence, medication exposure and progressive loss of function. Examine cognition and neurological status rather than relying on the mental-state examination alone. Reassessment matters because an evolving process may reveal itself through change rather than through a dramatic finding at the initial encounter.

The differential can name chronic subdural collection, brain tumour, raised intracranial pressure, hydrocephalus, vascular cognitive impairment and seizure-related presentations without reproducing their lesion associations or diagnostic details. Rank by urgency and plausibility. An immediately dangerous possibility may lead the action plan even when another diagnosis leads the probability list.

Investigation should be question led. State what must be detected, what modality is realistically available, and what escalation is needed if local capability is insufficient. An abnormal result must still be reconciled with the clinical pattern, while an unrevealing early result does not cancel a worsening trajectory. In Indian services, pragmatic sequencing may be necessary, but cost or access should not be allowed to convert a neurological warning pattern into routine psychiatric follow-up.

Where a broader neurocognitive syndrome is suspected, cross-refer the general diagnostic workup to the Dementias and neurocognitive disorders notes. Where episodic phenomena raise seizure questions, classification and anticonvulsant pharmacology remain with the Epilepsy and psychiatry notes.

Serial assessment is part of investigation. Document the baseline against which change will be judged, repeat the relevant mental-state and neurological observations, and communicate explicit triggers for escalation. A patient whose trajectory worsens despite an initially plausible psychiatric formulation needs the formulation reopened, not merely a stronger version of the same symptomatic treatment.

30.8 Convert the solution into an examination answer

A defensible answer has visible architecture. Open with a problem representation and provisional syndrome, then state the urgent alternative. Explain which observations support the formulation, which findings are missing, and how bedside examination and investigation will discriminate among the leading possibilities. End with management and reassessment rather than with the diagnostic label.

For management, specify **locus**, **focus** and **modus**. Locus identifies whether outpatient care, admission or emergency escalation is appropriate. Focus identifies the immediate target, the near-term functional goals and the continuing risks. Modus integrates cause-directed medical care, cautious symptom treatment, environmental and nursing measures, psychological support, rehabilitation, family work and social planning. Refer to current guidance for disease-specific treatment rather than naming an unverified guideline.

- State the leading formulation in qualified language.
- Name the dangerous alternative and the reversible contributor.
- Link examination and investigation to explicit clinical questions.
- Describe disposition, immediate safety and multidisciplinary care.
- Define what will be monitored and what would trigger reformulation.

MNEMONIC

TRACE the vignette

Tempo and trigger frame the change; **R**isk decides urgency; **A**rousal and attention test whether the interview is interpretable; **C**ognitive and neurological convergence supports localisation; **E**ffect on function, environment and care determines the plan.

The mnemonic is a reasoning spine, not a substitute for disease knowledge. It prevents premature closure while allowing a concise answer. If evidence conflicts, say so and specify the next observation that would resolve the conflict. If more than one process is plausible, distinguish the primary formulation from contributors instead of forcing a false choice.

In every neuropsychiatric vignette, establish state and safety before converting behaviour into diagnosis or anatomy.

31 · Consolidation and quick review

The marks lie in disciplined synthesis. A good answer does not recite every neuropsychiatric syndrome. It converts an apparently psychiatric presentation into a defensible clinical sequence: describe the change, establish its tempo, identify danger, examine the cognitive and neurological pattern, localise cautiously, construct the aetiological differential, and plan care. This approach also prevents the opposite error, in which every unusual behaviour is treated as proof of structural disease.

Neurological psychiatry is difficult because the same complaint can arise from network dysfunction, a systemic disturbance, medication or substance effects, the psychological response to illness, or a primary psychiatric disorder. The candidate must therefore hold psychiatric phenomenology and neurological reasoning together. The aim of consolidation is not diagnostic bravado. It is to show how uncertainty becomes narrower, safer and more clinically useful as history, collateral information, examination and investigation are integrated.

31.1 Begin with tempo, trigger and baseline

Problem representation is the shortest accurate statement of what changed, in whom, over what kind of time course, and in what clinical setting. It should include the dominant syndrome without prematurely naming a disease. “New behavioural change with impaired attention after a medical event” is more useful than “psychosis,” because it preserves the temporal and cognitive information that guides the next question.

Tempo is the opening discriminator. Abrupt change raises a different group of possibilities from fluctuation, a stepwise course, a gradual progression, or a stable deficit following a defined injury. The sequence of symptoms matters as much as their presence. A psychiatric symptom that follows altered alertness, a focal neurological change or a seizure-like event is framed differently from a longstanding psychiatric syndrome later complicated by neurological illness. Always establish the previous level of cognition, personality, function and independence through collateral history.

Define the last reliably known baseline rather than accepting “normal before” as sufficient history. Place psychiatric, cognitive, neurological and systemic changes on the same timeline. Ask what immediately preceded the change, including illness, injury, procedure, medication alteration, intoxication, withdrawal and psychosocial stress. Separate persistent deficits from episodic phenomena and from symptoms that vary with setting or physiological state.

The timeline should expose competing causal stories. A symptom may be a direct expression of brain dysfunction, an indirect consequence of disability, a treatment effect, or an independent comorbidity. More than one pathway may be active. A concise answer acknowledges this without becoming vague: state the leading formulation, the dangerous alternative, and the observation that would distinguish them.

Baseline is multidimensional. Ask separately about occupational and domestic function, self-care, navigation, social judgement, emotional regulation, communication and use of medicines or finances. A patient may have preserved routine conversation while family members have quietly

taken over complex tasks. Conversely, longstanding dependence should not be mislabelled as recent decline. The relevant comparison is the patient with their own prior function, not with an idealised population norm.

31.2 Build parallel syndrome, site, cause and consequence formulations

A robust case formulation runs on parallel tracks. The **syndrome formulation** describes the clinical pattern: disturbance of alertness and attention, an amnestic presentation, executive-behavioural change, focal higher-function loss, affective symptoms, psychotic symptoms, or a mixed picture. The **anatomical formulation** asks whether the pattern suggests diffuse dysfunction, a lateralised deficit, a network disconnection or pressure-related compromise. The **aetiological formulation** orders plausible causes by urgency, likelihood and reversibility. The consequence formulation covers risk, disability, decision-making ability, caregiver burden and rehabilitation needs.

Syndrome	Site	Cause	Impact
DESCRIBE	LOCALISE	EXPLAIN	PLAN

These formulations constrain each other. A proposed site must explain the observed pattern rather than merely match one dramatic sign. A proposed cause must fit the tempo and context. A treatment plan must address the actual consequence, not only the diagnostic label. When the tracks disagree, do not force agreement. Recheck the history, the level of arousal, sensory and motor limitations, language, medication exposure and collateral information.

RULE **Describe before you localise, and localise before you explain.** Premature disease naming discards the very observations that should test the diagnosis.

This parallel method is especially useful in long answers. It allows the candidate to move from phenomenology to neuroanatomy without claiming that a psychiatric symptom has a unique lesion correlate. It also leaves room for systemic illness and psychological adaptation, which may shape presentation even when a neurological lesion is established.

31.3 Place safety before diagnostic elegance

The early clinical task is not to produce a refined localisation. It is to decide whether the patient may have an evolving neurological or systemic emergency. Reduced or fluctuating alertness, rapid deterioration, a new focal deficit, a severe new headache pattern, repeated vomiting, seizure-like events, recent head injury, fever, marked physiological disturbance or signs suggesting pressure effects should move the assessment toward urgent medical care. These findings are signals for escalation, not diagnoses by themselves.

Diagnostic overshadowing occurs when a psychiatric label causes clinicians to discount new neurological or medical information. The reverse form is also common: an identified lesion is made responsible for every symptom, obscuring medication effects, pain, sleep disruption, demoralisation or an independent psychiatric disorder. Safety improves when the assessor asks paired questions explicitly: “What cannot be missed?” and “What else could be contributing?”

Clarify airway, breathing, circulation, physiological stability and level of consciousness before an elaborate psychiatric interview. Assess immediate risk arising from confusion, impulsivity, aggression, self-neglect, falls, wandering or inability to cooperate with essential care. Seek collateral information when the history may be unreliable because of impaired attention, memory, language, awareness or judgement. Escalate when the trajectory is worsening even if an earlier examination or investigation was reassuring.

PITFALL

A normal-looking conversation can conceal major cognitive or focal dysfunction. Social fluency, humour and confident answers do not establish intact attention, memory, executive control, awareness or neurological status. Test the relevant domains and verify the history.

Urgency depends on trajectory and vulnerability as well as symptom intensity. Quiet withdrawal may be more medically concerning than conspicuous agitation. Similarly, a subtle change noticed by family may carry more weight than a polished interview when it represents a clear departure from baseline. The safe answer explains why a patient can look psychiatrically settled yet still require urgent reassessment.

31.4 Use the bedside examination to test hypotheses

The bedside assessment should be hypothesis driven but ordered. Begin by establishing arousal and capacity to engage, then examine attention, orientation, language, memory, executive control, visuospatial function, praxis, recognition, neglect and relevant neurological signs. The purpose is not to accumulate abnormal findings. It is to determine whether the deficits form a coherent pattern and whether that pattern remains stable when task demands change.

State dependence means that performance may reflect fatigue, pain, intoxication, withdrawal, anxiety, sensory impairment, poor comprehension or fluctuating alertness rather than a fixed cognitive deficit. Before interpreting an error, confirm that the patient heard, understood and could physically perform the task. Repeat or simplify instructions without coaching the answer. Compare free conversation with structured testing and compare current performance with reliable collateral history.

A localising hypothesis becomes credible when several converging observations support it. A single bedside error is rarely sufficient. Language disturbance may contaminate memory and executive tasks; motor weakness may invalidate praxis testing; visual loss may mimic neglect or recognition failure. The examination therefore needs internal controls: use another modality, reduce motor demands, clarify comprehension, and observe spontaneous behaviour outside the formal task.

GOOD TO REMEMBER

When a patient “fails” a cognitive task, ask what the task actually required. Attention, language, vision, motor output and motivation may each sit between the target ability and the observed answer.

Record behaviour descriptively. “Did not search the affected side despite cueing” is more useful than an unsupported lobe label. Standardised cognitive instruments can support communication and follow-up, but their content, scoring and interpretation belong to the Psychotherapies, clinical assessment, MSE and nosology notes. The full dementia syndrome and its diagnostic workup belong to the Dementias and neurocognitive disorders notes.

31.5 Organise the differential by pattern rather than by memorised list

A neurological psychiatric differential becomes manageable when it is built from the observed pattern. Begin with tempo and consciousness, add focality and dominant cognitive domain, then incorporate the precipitating context. Finally, ask whether medication, substance exposure, systemic illness or psychosocial response modifies the picture. This produces a ranked differential rather than a catalogue.

PATTERN FRAME	QUESTION IT RAISES	COMMON REASONING ERROR
Acute or fluctuating global change	Is a diffuse medical, toxic or physiological disturbance active?	Calling it a primary psychiatric syndrome before establishing attention and arousal
Stable focal higher-function deficit	Do language, spatial, praxis, recognition or executive findings converge anatomically?	Localising from an isolated failed task
Predominantly amnestic presentation	Is encoding, storage, retrieval, temporal context or monitoring most affected?	Treating every inaccurate account as deliberate or psychotic
Behavioural or personality change	Is there executive-network dysfunction, an adaptive response, medication effect or psychiatric comorbidity?	Equating disinhibition or apathy with a single lesion
Progressive psychiatric-cognitive change	Could a structural, vascular or degenerative process be evolving?	Waiting for an obvious neurological deficit before reconsidering the formulation
Symptoms after injury or vascular event	Which features reflect direct injury, secondary complications, disability or independent illness?	Attributing the entire course to the original event

The table is a reasoning aid, not a diagnostic rule. Mixed presentations are expected because brain disorders affect networks, and because the person’s prior vulnerabilities, social context and treatment exposures shape expression. A ranked differential should explain the leading possibility, identify a dangerous alternative, name a reversible contributor, and state what information would change the ranking.

A useful ranking also separates probability from priority. A common explanation may lead the probability list, while a less likely but rapidly harmful condition leads the action list. State each explicitly. This prevents “most likely” from becoming “only possibility,” and it shows why urgent investigation can be justified without overstating diagnostic certainty. Re-ranking should follow

new collateral evidence, serial examination and treatment response, not attachment to the opening impression.

Specific syndrome criteria, lesion associations and management evidence belong in their dedicated sections. Substance-withdrawal syndromes are covered in the Substance use disorders notes. Seizure classification and anticonvulsant pharmacology are covered in the Epilepsy and psychiatry notes. The purpose here is to preserve the logic linking those bodies of knowledge.

31.6 Investigate to answer questions, not to display breadth

Question-led investigation begins by stating what must be detected or excluded. The urgent layer looks for instability and conditions in which delay could cause harm. The core layer searches for common medical, toxic, medication-related and structural explanations suggested by the history and examination. The focused layer follows the leading hypothesis with tests chosen for a specific discriminating purpose. A test should enter the answer because its result could change diagnosis, urgency, treatment or disposition.

Investigation never substitutes for bedside reassessment. A changing patient may need the formulation revised even when an earlier test was unrevealing. Conversely, an abnormal result may be incidental and should not displace a better clinical explanation without concordance of tempo, pattern and anatomy. State what each investigation is expected to clarify and how discordant findings will be handled.

Review prescribed, over-the-counter and traditional medicines, recent additions, dose changes, adherence and access. Ask directly about alcohol and other substances, including intoxication, withdrawal and interactions. Use collateral records to reconstruct prior cognition, function, injuries, vascular events, seizures and previous imaging or laboratory abnormalities. Reconcile each test result with the clinical pattern before assigning causality.

Pragmatism matters in Indian services, but it should be expressed as sequencing rather than lowered standards. Use the fastest locally available route that safely answers the immediate question, arrange transfer when the required capability is unavailable, and document why a deferred investigation is safe. Cost and access influence order, not the obligation to recognise danger.

Detailed imaging patterns, screening operations, treatment regimens and disease-specific investigations remain with their owning sections. Emergency nutritional syndromes and repletion protocols are covered in the Endocrine and metabolic psychiatry notes. The consolidating skill is to make every requested test traceable to a clinical question.

31.7 Manage by locus, focus and modus

Management begins with **locus**: decide whether care is safe in the outpatient setting, requires admission, or needs emergency or higher-dependency support. The decision follows physiological stability, trajectory, behavioural risk, supervision, capacity to cooperate, caregiver reliability and access to reassessment. A diagnostic label alone does not determine setting.

Focus identifies what must change now and what can wait. Acute priorities are stabilisation, treatment of the underlying cause, prevention of secondary harm and control of behaviour only when it threatens safety or essential care. Short-term work includes sleep, pain, communication, mobilisation, environmental support and early rehabilitation. Longer-term care addresses residual cognitive, emotional and behavioural problems, recurrence prevention, disability, family education, vocational goals and caregiver strain.

Modus ensures that treatment is not reduced to prescribing. Cause-directed medical care, cautious symptom treatment, psychological support, rehabilitation, environmental modification, nursing strategies, family work and social intervention should be considered together. Choose the least restrictive effective approach and review frequently because brain-injured or medically unwell patients may be unusually vulnerable to sedation, cognitive worsening, falls, interaction effects and lowered participation in rehabilitation.

TRAP **Do not present a drug list as a management plan.** Examiners look for setting, target, cause-directed care, non-pharmacological measures, monitoring, rehabilitation and follow-up as well as symptom control.

Drug pharmacology should be cross-referred rather than duplicated. Antidepressant selection belongs to the Mood disorders: pharmacotherapy notes. Seizure drug choice belongs to the Epilepsy and psychiatry notes. Withdrawal regimens belong to the Substance use disorders notes. When current guidance is needed, consult the relevant up-to-date guideline and local protocol; do not convert a remembered regimen into a universal rule.

Monitoring should be tied to the anticipated benefit and harm. Define the target behaviour or function in observable language, establish who will review it, and decide what deterioration triggers escalation. Involve family without transferring clinical responsibility to them. Explain expected uncertainty, warning changes, environmental strategies and the purpose of follow-up. Rehabilitation begins during acute care through attention to communication, mobility, sleep-wake organisation and participation, then evolves with the person's recovery and social role.

31.8 The TIERED master method for cases and essays

Use one reusable spine under pressure. The master mnemonic is **TIERED**. It does not replace syndrome knowledge. It tells the candidate in what order to deploy that knowledge, whether the stem concerns a vascular event, head injury, acute cognitive change, amnesia, a focal higher-function disorder, a structural lesion or a seizure-related presentation.

MNEMONIC

TIERED

Tempo and trigger: define baseline, onset, sequence and course. **I**mmediate threat: identify instability, deterioration and risks that change the care setting. **E**xamination pattern: describe arousal, cognition, behaviour and neurological findings. **R**egional hypothesis: localise only when findings converge. **E**tiology and effects: rank causes and state functional, behavioural and caregiver consequences. **D**isposition and domains of care: apply locus, focus and modus, then specify monitoring and follow-up.

For a short case, compress the spine into a problem representation, leading diagnosis, dangerous alternative, discriminating findings and immediate plan. For a long essay, expand each letter into a paragraph and insert disease-specific evidence from the relevant section. In either format, keep observation separate from inference. State “the pattern suggests” when localisation is probabilistic, and reserve definitive language for findings that genuinely establish the conclusion.

- Name the dominant syndrome and the main competing formulation.
- Show why tempo, trigger and examination pattern favour one over the other.
- State the urgent risk and the reversible contributor that require active exclusion.
- Close with setting, cause-directed treatment, symptom targets, rehabilitation and follow-up.

The enduring habit is calibrated reasoning. Psychiatric symptoms can be neurologically meaningful without being anatomically unique. Neurological disease can be present without explaining the whole person. Collateral history can outweigh a fluent interview, and a changing trajectory can outweigh a reassuring earlier snapshot. The best answer makes these tensions explicit and resolves them through structured reassessment rather than assertion.

TIERED: establish tempo, detect immediate threat, define the examination pattern, localise cautiously, rank aetiology and effects, then decide disposition and the full domains of care.

Sources and further reading

The clinical specifics in this chapter are grounded in the standard texts below; where two sources set different thresholds, both are named in the text as a conflict rather than reconciled. Instrument content (the cognitive screens) is described and cross-referred, never reproduced.

Sadock BJ, Sadock VA, Ruiz P.
Kaplan and Sadock's Comprehensive Textbook of Psychiatry. 11th ed. Wolters Kluwer, 2024. The primary reference for the neuropsychiatric syndromes, delirium and the amnestic disorders.

Arciniegas DB, Yudofsky SC, Hales RE (Kaufman).
Kaufman's Clinical Neurology for Psychiatrists. 8th ed. Elsevier. Focal syndromes, traumatic brain injury and the bedside examination.

Brazis PW, Masdeu JC, Biller J.

Localization in Clinical Neurology. 8th ed. Wolters Kluwer. Lobe and vascular localisation.

Tasman A, Kay J, Lieberman JA, et al.
Tasman's Psychiatry. 5th ed. Springer. Post-stroke and post-injury neuropsychiatry.

Geddes JR, Andreasen NC, Goodwin GM (New Oxford).
New Oxford Textbook of Psychiatry. 3rd ed. Oxford University Press. Vascular cognitive impairment.

Johnstone EC, Owens DGC, Lawrie SM, et al.
Companion to Psychiatric Studies. 8th ed. Elsevier. Vascular criteria and inherited small-vessel disease.

American Psychiatric Association.
Diagnostic and Statistical Manual of Mental Disorders, 5th ed, Text Revision (DSM-5-TR). APA, 2022; and the World Health Organization ICD-11 Clinical Descriptions and Diagnostic Requirements (CDDR).

Taylor D, Barnes TRE, Young AH.
The Maudsley Prescribing Guidelines in Psychiatry. 14th ed. Wiley-Blackwell; and Stahl SM, Prescriber's Guide, Stahl's Essential Psychopharmacology. Cambridge University Press. Prescribing in acquired brain injury and pseudobulbar affect.

Hodges JR.
Cognitive Assessment for Clinicians. 3rd ed. Oxford University Press; and Kipps CM, James J. Cognitive Neurological Examination. The bedside cognitive method and its localising signs.

PART VI

VI

Neuropsychiatry of Systemic and Neurological Disease

The systemic and neurological diseases that reach psychiatry through inflammation, infection and degeneration, and how each announces itself.

01 HIV and its neuropsychiatric syndromes

02 Parkinson disease and its neuropsychiatry

03 Huntington disease and other movement disorders

04 Autoimmune and infective encephalitides

05 Neurosyphilis and CNS infections

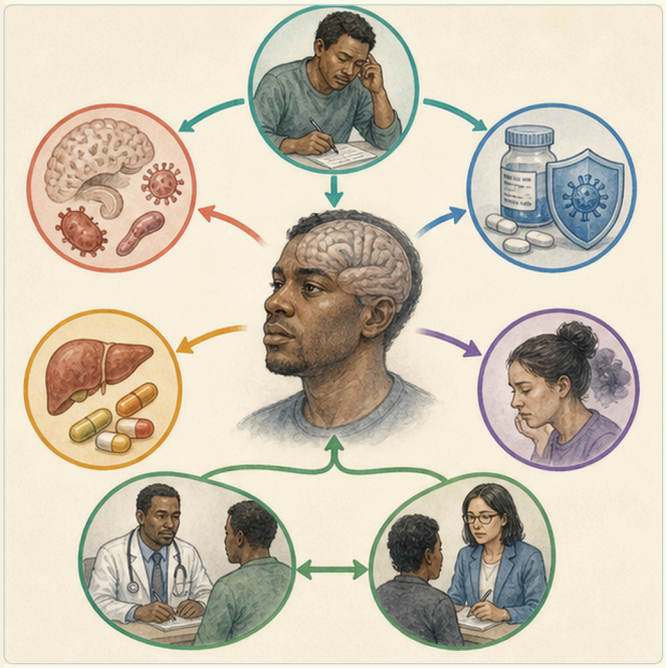
06 Multiple sclerosis and psychoneuroimmunology

07 Degenerative, neuromuscular and functional neurology

08 Consolidate and apply

09 Sources and further reading

VISUAL 6.1



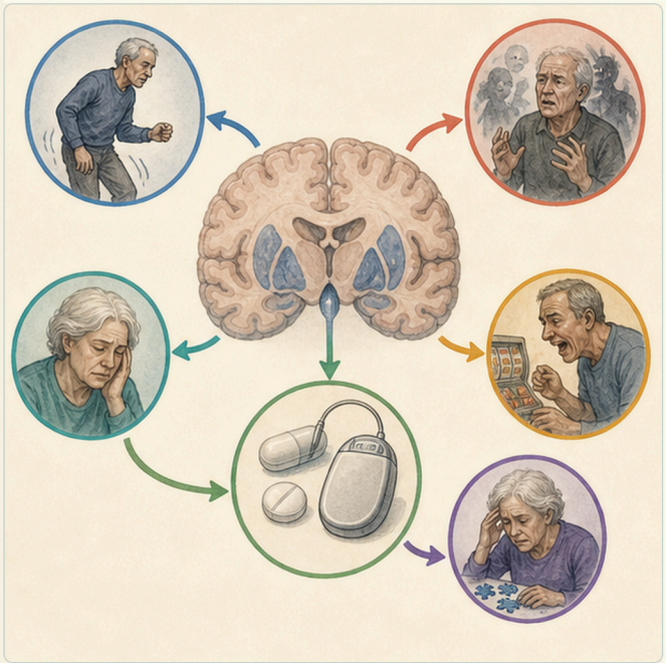
HIV neuropsychiatry is a layered differential

- HAND
- Opportunistic infection
- ART and viral control
- Mood and psychosis
- Drug interactions
- Continuity of care

Neuropsychiatric symptoms in HIV require separation of HAND, infection, treatment effects, primary psychiatric illness and drug interactions.

Source lock: day-35.final.html#sec-t06

VISUAL 6.2



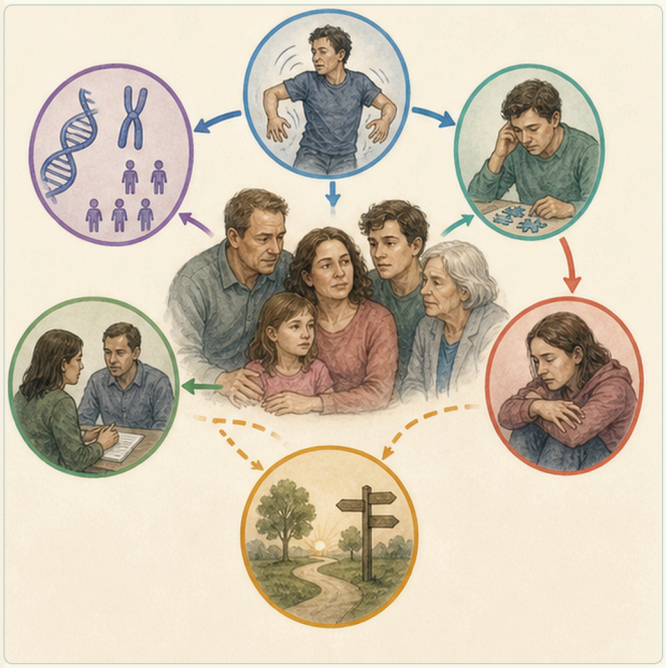
Parkinson treatment changes psychiatric risk

- Motor state
- Psychosis
- Mood and anxiety
- Impulse control
- Cognition
- Treatment review

Parkinson neuropsychiatry is shaped by disease, cognition and dopaminergic or device treatment, so every intervention is reviewed across motor and psychiatric outcomes.

Source lock: day-35.final.html#sec-t11

VISUAL 6.3



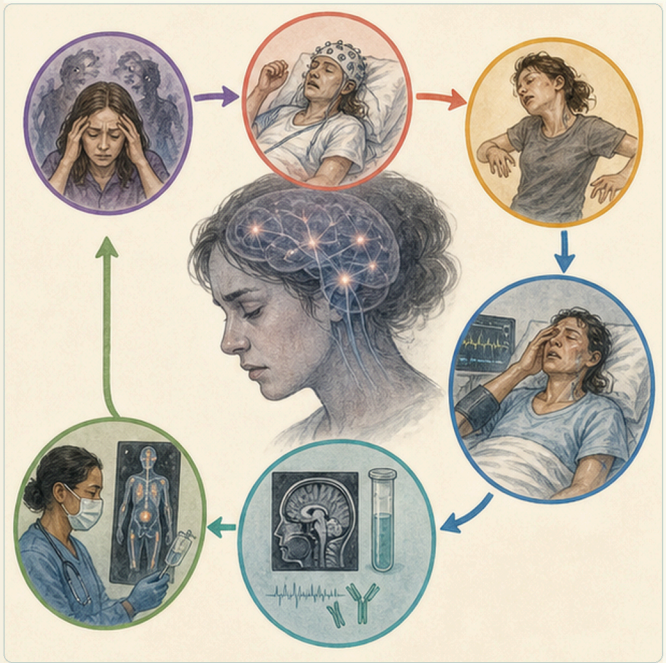
Huntington disease joins genetics to lived decisions

- HTT inheritance
- Movement
- Cognition
- Psychiatric change
- Predictive counselling
- Future planning

Huntington disease links movement, cognition and psychiatric change to an autosomal-dominant diagnosis whose predictive testing requires counselling.

Source lock: day-35.final.html#sec-t13

VISUAL 6.4



Autoimmune encephalitis crosses psychiatric boundaries

- Psychiatric change
- Seizures and catatonia
- Abnormal movements
- Autonomic instability
- CSF EEG MRI antibodies
- Tumour search and immunotherapy

Autoimmune encephalitis is suspected when psychiatric presentation is joined by neurological, autonomic or investigation red flags.

Source lock: day-35.final.html#sec-t18

PAPER IV

Neuropsychiatry of systemic and neurological disease

Seven bands . one complete notes document

HIV and its neuropsychiatric syndromes

1 · Systemic and neurological disease for the psychiatrist: what this chapter owns, and what it points to

A psychiatric presentation is sometimes the first visible sign of a disease that lives in the body or the nervous system rather than in the mind alone. A patient may reach the clinic with depression, mania, psychosis, anxiety, apathy, personality change or failing memory while the process driving the presentation is an infection, an autoimmune attack, a degenerative disorder, a demyelinating disease or the treatment given for one of these. The mental state names what the patient shows. It does not, on its own, name the disease that produced it.

This orientation supplies the **shared method** for the chapter. The dedicated sections teach each disease, its neuropsychiatric features, its investigations and its treatment in full. The task here is different and more portable: to learn how to recognise when a psychiatric syndrome is secondary to a systemic or neurological illness, how to reason from surface to mechanism, how far to pursue the underlying disease, and how to treat the illness and the mind together. This is where the marks sit, because a memorised list of disease-and-symptom associations is not yet a formulation.

1.1 Define the territory before naming the disorder

The **neuropsychiatry of systemic and neurological disease** concerns changes in mood, thought, perception, behaviour, motivation and cognition that arise as manifestations of a definable bodily or brain disease. The unifying idea is that the psychiatric syndrome is the surface and a diagnosable process, infective, immune, degenerative, structural, metabolic or iatrogenic, is the depth. The central difficulty is that the surface looks like ordinary psychiatry. Low mood, elation, delusions, hallucinations, obsessional phenomena, apathy and forgetfulness can each be produced by disease of the brain or the body just as they are produced by primary psychiatric illness.

The opening question is therefore not only which psychiatric diagnosis this is, but what could be producing this presentation, and whether the story behaves like a primary disorder or a secondary one. A **secondary neuropsychiatric syndrome** earns its label from context: an atypical age at onset, an atypical mixture of symptoms, accompanying physical or neurological signs, a course that does not run as the primary illness runs, or a systemic disease already known. Holding the presentation as new psychiatric change in the setting of a medical or neurological illness preserves the sequence and the uncertainty that a premature primary label would throw away.

The same discipline guards against **diagnostic overshadowing** in both directions. A known psychiatric history should not make a new organic change invisible, and a demonstrable disease should not be made to explain every symptom merely because it is concrete. Fear, pain, insomnia, loss, family strain, the effects of treatment and premorbid vulnerability all remain clinically active even when a systemic or neurological disease is established.

RULE **A new or atypical psychiatric syndrome arising in the setting of systemic or neurological disease should be treated as potentially secondary until the disease has been assessed.** Describe the mental state first, ask what disease could produce it, then treat the disease and the syndrome together rather than choosing between them.

1.2 Read the chapter as a map of diseases

The chapter is arranged by the kind of disease confronting the psychiatrist. These are not sealed compartments: an immune process may follow an infection, a degenerative disease may open with a mood disorder, and the treatment of one disease may itself disturb the mind. The value of the map is to show which lens leads and where the detailed teaching lives, so that an answer stays centred on the disease actually in front of the patient. The accompanying figure lays the same territories out as a single reference frame.

DISEASE TERRITORY	OPENING QUESTION	DEPTH SUPPLIED BY THE CHAPTER
HIV and its neuropsychiatric burden	Is this cognitive, affective or psychotic change a direct effect of the virus, an opportunistic complication or a drug effect?	The cognitive spectrum, the neuro-injury and its response to antiretroviral therapy, psychiatric comorbidity, opportunistic infections and drug interactions
Parkinson disease and the parkinsonian disorders	Are these symptoms the disease, its treatment or an independent illness?	Psychosis and antipsychotic choice, depression and anxiety, impulse and dopaminergic behaviour, cognition, stimulation effects and the atypical parkinsonian mimics
Huntington disease and the hyperkinetic disorders	How does an inherited or acquired movement disorder shape mind and behaviour?	The genetics and its counselling, the psychiatric and cognitive course, and the chorea of the other movement disorders
Autoimmune, limbic and viral encephalitides	Could an antibody or a virus be producing this subacute psychiatric change?	The antibody-mediated encephalitides, the red flags for an autoimmune psychosis and the viral encephalitides with psychiatric sequelae
Neurosyphilis and the chronic CNS infections	Is a treatable chronic infection presenting as a psychiatric or cognitive disorder?	Neurosyphilis and general paresis, the endemic neuroinfections and the subacute and rare chronic infections
Multiple sclerosis, demyelination and immunity	How do a relapsing brain disease and the immune system reach the mind?	The psychiatric manifestations of demyelination, the effect of disease-modifying treatment and the immune hypothesis of psychiatric disorders
Degenerative, neuromuscular and functional disorders	Is this a degenerative, neuromuscular or functional process, and can its positive signs be recognised?	Motor neuron disease, the neuromuscular and mitochondrial disorders and the functional neurological disorder spectrum

Disease	**Route**	**Syndrome**	**Care**
WHAT IS WRONG	HOW IT REACHES THE BRAIN	WHAT THE PATIENT SHOWS	WHAT TREATMENT MUST DO

Read the map from either direction. A known systemic or neurological disease directs the clinician toward the psychiatric and cognitive syndromes it can produce. An unexplained psychiatric change directs the clinician back toward the diseases that must not be missed. The sections supply the depth; the orientation preserves integration, so that disease, mechanism, syndrome and care remain part of one formulation rather than four disconnected recollections.

TRAP **Reciting disease-and-symptom associations instead of reasoning through the case.** Candidates lose marks by producing a catalogue of which disease causes which symptom. The examiner is testing a method: recognise the features that mark a syndrome as secondary, identify the disease, investigate it and treat both fronts. A list demonstrates recall; the method demonstrates a clinician.

1.3 Reason from surface to mechanism

Depth in this chapter means asking how a disease reaches the mind. A handful of **mechanistic routes** recur across every territory, and naming the route that applies both explains the presentation and points toward the treatment. The routes are not mutually exclusive, and more than one may be active in the same patient, but each answers a different question about how the surface syndrome was produced.

Direct involvement is the disease acting on the brain itself, whether by infection of neural tissue, by structural injury or by a genetically driven degeneration of specific circuits. **Immune and inflammatory injury** is antibody-mediated, inflammatory or demyelinating disturbance of neural function, sometimes arising remotely from any active infection. **Infective and post-infective** mechanisms cover the organism itself and the host response to it, producing an acute encephalitic picture or a chronic smouldering one. **Systemic and metabolic** routes are disease elsewhere in the body altering the internal environment on which the brain depends. **Iatrogenic** effects are the neuropsychiatric consequences of the treatment given for the disease, or for the psychiatric syndrome itself.

The **iatrogenic neuropsychiatric syndrome** deserves particular vigilance, because it is common, often reversible and frequently missed. The drug that treats the disease and the drug that treats the mind can each disturb cognition, mood, sleep and behaviour, and they can interact so that the effect is greater than either alone. A careful drug history, including the timing of any change against the timing of the syndrome, is therefore a diagnostic instrument rather than a formality. Reaching for the mechanism early is what converts a description into a plan, because the route named is usually the route that treatment must interrupt.

1.4 Recognise organicity through red flags

Because secondary syndromes imitate primary ones, recognition rests on **red flags** rather than on the psychiatric phenotype alone. No single flag is decisive, but their convergence should redirect the assessment from a primary diagnosis toward a search for disease. The recurring flags are an onset at an age unusual for the apparent primary disorder, a mixture of symptoms that does not fit a familiar syndrome, coincident physical or neurological signs, disturbance or fluctuation of consciousness, cognitive decline out of proportion to the affective picture, seizures or other paroxysmal events, systemic features such as fever or weight loss, and a course that is unusually rapid or resistant to appropriate psychiatric treatment.

Organicity becomes probable when several of these run together, and it becomes more urgent when the tempo is subacute, when consciousness is involved or when the trajectory is worsening despite treatment. The reasoning is probabilistic rather than mechanical: the flags raise the pre-

test probability of a secondary syndrome and lower the threshold for examination and investigation. They do not, by themselves, name the disease, and a patient with none of them may still harbour one, so their absence reassures less than their presence alarms.

PITFALL

Formulating without examining. A psychiatrist who omits fundoscopy, a basic neurological screen and a general physical examination can miss papilloedema, a focal sign or a systemic clue that would have redirected the entire diagnosis. The organic sign is found only by the clinician who looks for it, and the examination that finds it takes minutes while the illness it reveals may take a life.

1.5 Investigate to answer a question, not to reassure

Investigation in this chapter is hypothesis-led, not reflexive. The purpose of **targeted investigation** is to answer the mechanistic question the history and examination have raised, not to order every available test in the hope that one will speak. Structural imaging answers a question about structure and mass. Examination of the cerebrospinal fluid answers questions about infection and inflammation. Serological and specific antibody testing answers immune and infective questions. Electrophysiology answers questions about cortical function and paroxysmal activity. Each modality has a question it settles well and questions it cannot settle at all, and choosing the test means first naming the question.

The workup is protected by discipline in sequence and in interpretation. Sequence means letting the leading hypothesis order the tests, so that the investigation most likely to change management is done first and the staged workup is completed rather than abandoned. Interpretation means remembering that a normal result rarely excludes a disease outright. A single unremarkable investigation early in a subacute illness can mislead, and serial testing, a different modality or a later repeat may be needed when clinical suspicion is strong. A result is evidence to be weighed against the whole picture, not a verdict that overrides it.

GOOD TO REMEMBER

In any first episode of psychosis, mania or severe depression with atypical features, examine the nervous system and take a full drug, infection and systemic history before committing to a primary diagnosis. The examination is quick and repeatable; the secondary disease it protects against is neither.

1.6 Know what this chapter owns and what it hands over

The chapter touches many constructs that are taught in full elsewhere. A disciplined answer names the interface and cross-refers; it does not reproduce another chapter's criteria, scoring table or pharmacology. Knowing the **diagnostic boundary** is itself part of competence, because it keeps the answer on the disease question actually asked rather than drifting into a neighbouring syndrome.

- Pseudobulbar affect, the pathological laughing and crying that appears in multiple sclerosis and in motor neuron disease, is named here in those diseases but taught in full, with its

phenomenology and treatment, in the Stroke, TBI, delirium and focal brain syndromes notes.

- Dementia with Lewy bodies and its temporal distinction from Parkinson disease dementia belong to the Dementias and neurocognitive disorders notes; the chapter teaches the Parkinson-disease side of that boundary and cross-refers for the rest.
- The diagnostic criteria and phenomenology of Tourette syndrome belong to the Child behavioural, emotional and other disorders notes; the chapter names it as the archetypal tic disorder while teaching the acquired choreas itself.
- Catatonia phenomenology, its structured rating and the lorazepam challenge belong to the Other psychotic disorders, delusional syndromes and catatonia notes; the chapter teaches only how an autoimmune encephalitis is distinguished from catatonia.
- The primary nosology and pharmacotherapy of depression and schizophrenia belong to the Mood disorders: pharmacotherapy notes and the Schizophrenia: management, antipsychotics and adverse effects notes; the chapter teaches only the secondary and immune-mediated versions of these syndromes.
- The hepatic and metabolic aspects of Wilson disease belong to the Endocrine and metabolic psychiatry notes, and the standardised cognitive instruments, with their item content and scoring, belong to the Psychotherapies, clinical assessment, MSE and nosology notes.

A useful handover sentence does real work: it names the neighbouring construct, says why it matters to the case in hand, and points to its owning chapter, then returns to the disease question. This is shorter and more accurate than inserting a copied criterion set, a scoring table or a drug-selection essay into every answer that happens to touch them.

1.7 Manage the disease and the mind together

Management follows the shared spine, which keeps an answer complete under pressure. The locus of care asks where treatment can safely happen: outpatient and community care for the stable patient, inpatient care when medical or behavioural need exceeds community capacity, and emergency care when an encephalitic, infective or metabolic crisis threatens life or brain. The focus orders the targets over time: acute work treats the underlying disease and the immediate risk, shorter-term work manages the disturbing psychiatric symptoms, and continuing care addresses residual cognitive, affective and behavioural change, relapse prevention and rehabilitation. The modus ensures every mode is considered: disease-directed medical treatment, cautious psychiatric treatment, psychological support, procedural options where the disease indicates them, and social and family intervention.

The two prescriptions, one for the disease and one for the mind, must be planned together, because each can worsen the other and because the choice and dose of a psychotropic often depend on the specific disease and its own treatment. Where a management decision rests on a formal guideline position, that position should be attributed to the body that issues it and drawn from current disease-specific guidance and local protocol rather than recalled as a rule of thumb. The reasoning is taught here; the reader is pointed to current guidance for each disease in its own section.

IN INDIA

Where neurocysticercosis, Japanese encephalitis, cerebral malaria, central nervous system tuberculosis, HIV and neurosyphilis remain part of everyday practice, keep these treatable diseases in the differential of a new psychiatric or cognitive presentation that carries neurological or systemic signs, and investigate the brain and cerebrospinal fluid as the suspected syndrome directs, weighing the contraindications to lumbar puncture, before settling on a primary psychiatric diagnosis. Treating a curable infection as a functional illness is the costliest error this chapter guards against.

1.8 Carry the ORGANIC scaffold into essays and vivas

Under examination pressure, the method must survive when detail fails. The master mnemonic for this chapter is **ORGANIC**. It is an educational scaffold rather than a diagnostic instrument, and it organises any case in the chapter, a viral or autoimmune encephalitis, a movement disorder, a chronic infection, a demyelinating disease or a treatment effect, without replacing the disease-specific knowledge that lives in the owning section.

MNEMONIC

ORGANIC

Onset and course: is the mental-state change atypical in age, tempo or mixture for a primary disorder? **R**ecognise the disease: examine the nervous system and the body for the illness behind the syndrome. **G**enerate the mechanism: direct, immune, infective, systemic-metabolic or iatrogenic. **A**ssess with targeted tests: image, sample the cerebrospinal fluid, serology and electrophysiology as the hypothesis demands. **N**ame the boundary: teach what this chapter owns and hand over what it does not. **I**ntervene on both fronts: treat the disease and the psychiatric syndrome together. **C**ontinue and coordinate: monitoring, rehabilitation and follow-up across settings and disciplines.

For a short case, compress ORGANIC into a problem representation, the leading disease, the dangerous alternative, the discriminating investigation and the immediate plan. For a long essay, expand each letter and insert the epidemiology, criteria, investigations and treatment from the owning section. In a viva, make the calibration audible: state what is observed, what is inferred, what remains uncertain, and which finding would change the ranking. The enduring habit is integration without overreach, so that a psychiatric syndrome can be recognised as the expression of a disease without the disease being made to explain the whole person.

ORGANIC: read onset and course, recognise the disease, generate the mechanism, assess with targeted tests, name the boundary, intervene on both fronts, then continue and coordinate care.

2 · HIV-associated neurocognitive disorders: the HAND spectrum and its classification

A person living with HIV who begins to slip cognitively has not necessarily acquired a dementia. Cognitive change is one of the commonest neuropsychiatric problems in HIV, but in

the treated era most of it is mild and much of it is silent, and the whole clinical value of the field lies in grading what you find rather than treating every positive test as a catastrophe. The marks concentrate here for two reasons. The first is a nomenclature that examiners love, because the disorders sit on a graded spectrum with a precise dividing line. The second is an epidemiological trap: two respectable sources give materially different category splits, and a candidate who quotes one as canonical has walked into it. This section sets out the three-tier **HIV-associated neurocognitive disorder** (HAND) spectrum, the framework that grades it, the clinical criteria that sit alongside the research nosology, and the prevalence figures with their populations kept firmly attached. The mechanism of the neuro-injury itself, and how antiretroviral therapy alters it, is taken up separately in this chapter and is only named here.

2.1 The three-tier HAND spectrum

Most authors divide HIV-associated neurocognitive impairment by severity into three groups, applied only when the presentation cannot be explained by another condition: **asymptomatic neurocognitive impairment** (ANI), **mild neurocognitive disorder** (MND), and **HIV-associated dementia** (HAD), set out in order of increasing severity. The ordering is the spine of the whole topic: ANI is the least severe and HAD the most, with MND between them. Holding that ladder in mind, and knowing what separates each rung from the next, is worth more in an examination than any single prevalence figure.

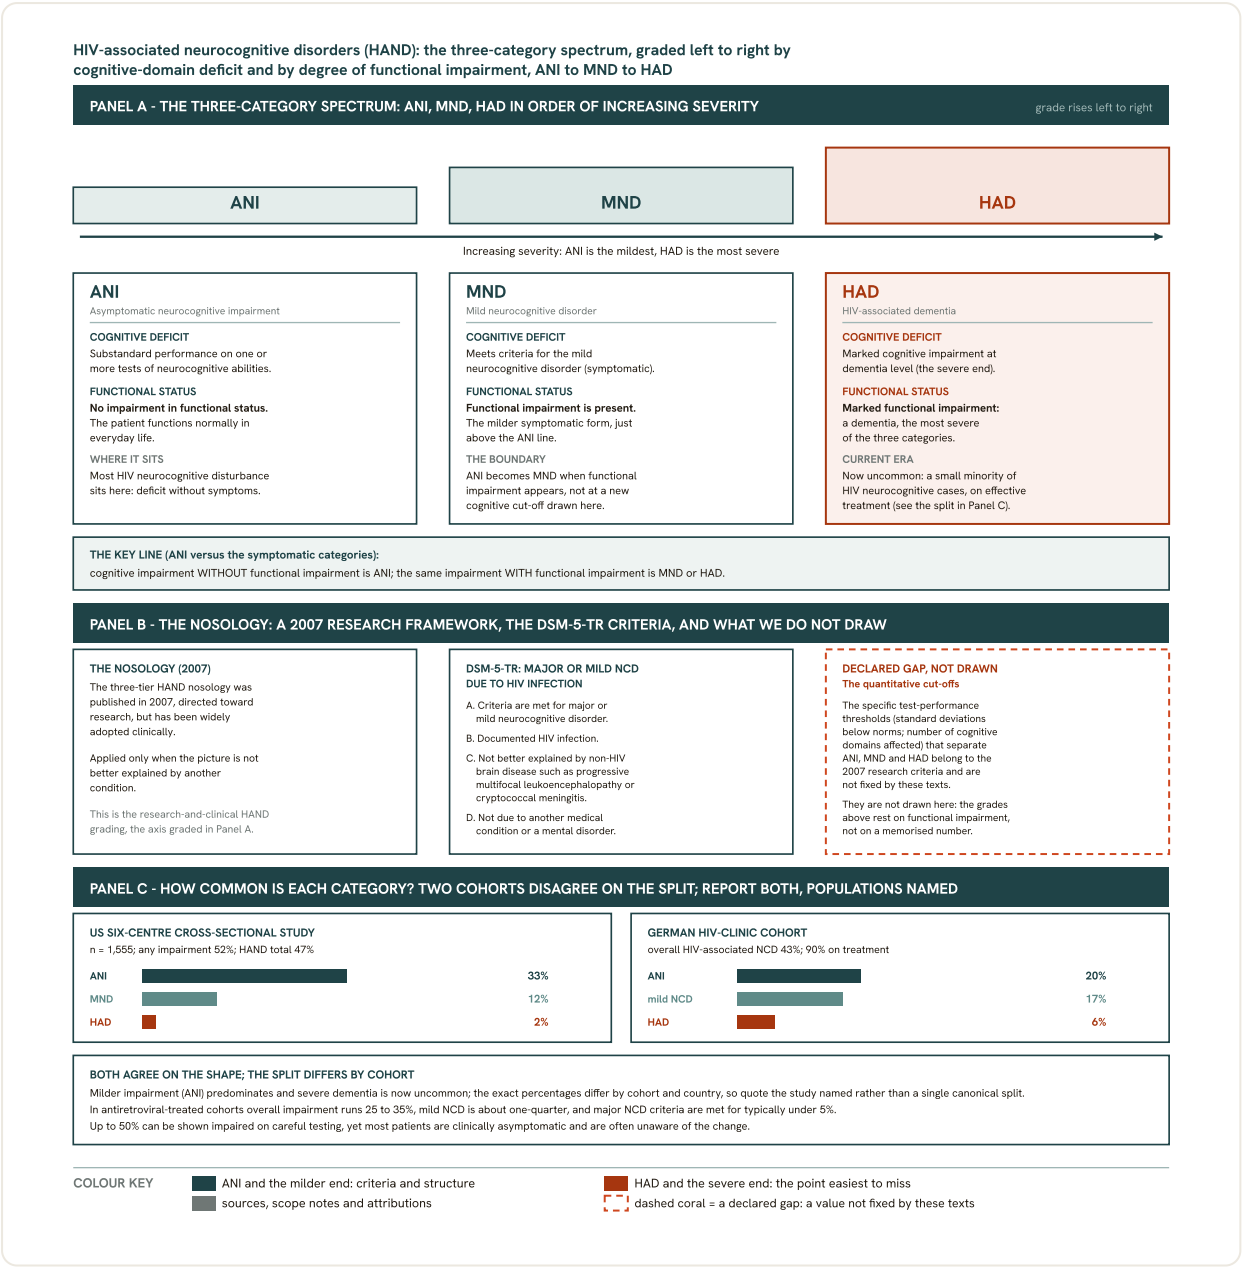


Fig 35.1 HIV-associated neurocognitive disorders (HAND): the three-category spectrum. Most authors grade HAND by increasing severity into asymptomatic neurocognitive impairment (ANI), mild neurocognitive disorder (MND) and HIV-associated dementia (HAD); the boundary from ANI to the symptomatic categories is the appearance of functional impairment, since ANI shows substandard test performance with intact everyday function. The 2007 research nosology is widely used clinically, but its quantitative test-performance cut-offs are not fixed by these texts and are not drawn here. Category rates differ by cohort: a US six-centre study (n = 1,555) found ANI 33%, MND 12% and HAD 2%, while a German clinic cohort found ANI 20%, mild NCD 17% and HAD 6%. (Kaplan and Sadock, CTP 2024; DSM-5-TR.)

The three tiers are best read as a single axis with two dividing lines rather than as three unrelated diagnoses. The lower line separates measured impairment that leaves the person functioning normally from impairment that has begun to bite into everyday life; the upper line separates that milder symptomatic disorder from frank dementia. The accompanying figure lays the spectrum out this way. What the tiers do not do is describe three different diseases: they are severity categories along a single axis rather than proven sequential stages of one process, which is why a patient can move between them over time and why re-grading at review matters more than the label first applied. Movement in either direction is possible, and this is a point of clinical

substance rather than pedantry: a patient can slip from asymptomatic to symptomatic impairment as function begins to fail, and can equally recover a tier when a reversible contributor is corrected. Because of that mobility, the tier is a snapshot of where the patient sits today, not a permanent classification, and treating it as fixed is one of the ways the grading is misused.

SEVERITY TIER	NEUROPSYCHOLOGICAL TESTING	EVERYDAY FUNCTIONAL STATUS
Asymptomatic neurocognitive impairment (ANI)	Substandard performance on one or more tests	No impairment in functional status
Mild neurocognitive disorder (MND)	Measured impairment present	Functional impairment present
HIV-associated dementia (HAD)	Most severe impairment of the three	Functional impairment present, at the level of dementia

2.2 The 2007 research nosology and the cut-offs it rests on

The currently used nosology was published in **2007** and was directed toward research, but it has been widely adopted clinically and has been the basis of significant discussion. That dual life, born as a research instrument and pressed into clinical service, explains much of the debate around it: a framework built to standardise study populations is being asked to label individual patients, and the fit is imperfect. Understanding that it is a **research nosology** first is the single most useful thing to know about where the categories come from.

The tiers are defined by the degree of measured neuropsychological impairment together with the presence or absence of functional decline. The precise psychometric thresholds, how many standard deviations below normative performance a patient must fall, and in how many cognitive domains, belong to the primary research criteria and are not reproduced here, because a standard reference that names the three categories and dates the nosology does not itself print those cut-offs. This is a real gap rather than an oversight: quoting a remembered standard-deviation rule as though it were grounded is exactly how a plausible but unverified number reaches a candidate. The safe teaching is that the cut-offs exist, that they are demographically anchored, and that they must be read from the primary criteria before being drawn as a grid.

RULE **HAND is a graded diagnosis of exclusion, and the first cut is functional, not psychometric.** The three tiers ascend ANI to MND to HAD, are applied only when the impairment cannot be explained by another condition, and the lowest tier is separated from the rest by whether everyday function is affected, not by the raw test score alone.

2.3 The functional-status hinge: ANI versus the symptomatic disorders

The distinction that carries the most weight, and the one most often tested, is the line between the asymptomatic tier and everything above it. Most neurocognitive disturbances in HIV do not meet the criteria for mild neurocognitive disorder and instead represent asymptomatic neurocognitive impairment: patients who show substandard performance on one or more tests of neurocognitive abilities but do not have any impairment in functional status. The pivot is **functional status**. Below the line, the deficit is detectable only on formal testing and the person

continues to work, manage medication and run a household unimpeded; at and above the line, the deficit has started to cost the patient something in daily life.

MNEMONIC

Impaired but intact, impaired and impaired, demented

The three HAND tiers in ascending severity. ANI: cognition impaired on testing but everyday function intact. MND: cognition impaired and everyday function now affected. HAD: HIV-associated dementia. The hinge between the first tier and the other two is functional status, not the test score.

Getting this hinge right changes both counting and management. It is why ANI, despite being defined by measured impairment, is called asymptomatic: the person does not experience or report a functional problem. It is also why ANI dominates the epidemiology, since the great majority of HIV-related cognitive impairment sits at this sub-threshold level. Whether ANI in an otherwise well patient is a benign finding or an early marker of trajectory remains genuinely contested, which is one more reason to grade rather than to dismiss.

2.4 The DSM-5-TR clinical criteria and what must be excluded

Running alongside the research nosology is the clinical diagnosis, **Major or Mild Neurocognitive Disorder Due to HIV Infection**, set out in DSM-5-TR. It is not a competitor to the ANI/MND/HAD scheme so much as its clinical counterpart: it fixes the diagnosis to a documented infection and, above all, insists that the deficit is not the footprint of something else. Its criteria run as follows:

- The criteria are met for major or mild neurocognitive disorder in the first place, establishing the level of the deficit.
- There is documented infection with human immunodeficiency virus.
- The disorder is not better explained by non-HIV conditions, including secondary brain diseases such as progressive multifocal leukoencephalopathy or cryptococcal meningitis.
- It is not attributable to another medical condition and is not better explained by a mental disorder.

The third criterion is where most clinical errors live. The CNS opportunistic infections it names are the classic secondary causes that mimic HAND, and they are taught in their own right elsewhere in this chapter; here they matter only as the conditions that must be off the table before the label is applied. HAND is, in the end, a diagnosis of exclusion resting on a documented infection.

The clinical diagnosis and the research tiers overlap but do not map cleanly onto one another, and it is worth holding that caution in mind rather than forcing a cell-for-cell translation. The major neurocognitive disorder end of the clinical diagnosis corresponds broadly to the dementia pole of the research spectrum. The milder end is where the two schemes diverge: DSM-5-TR mild neurocognitive disorder does not itself require a loss of independence, so it does not line up one-

to-one with the functional cut that separates asymptomatic impairment from the symptomatic research tiers. Seen this way, the DSM-5-TR diagnosis and the ANI/MND/HAD ladder are overlapping but non-equivalent views of one graded process, and the shared vocabulary of impairment and function is what keeps the two frameworks talking to each other.

PITFALL

Labelling new cognitive decline in a patient with HIV as HAND without first excluding the secondary brain diseases. An opportunistic CNS infection can present as subacute cognitive change, and DSM-5-TR makes the exclusion an explicit criterion precisely because the treatable mimic is easy to miss when the HIV diagnosis is already known and anchors the clinician's thinking.

2.5 How common HAND is, and why most of it is quiet

Cognitive disturbance in HIV is common by any measure. Between about one-third and over one-half of HIV-infected individuals have at least some evidence of a neurocognitive disturbance, and on formal assessment the figure reaches the upper end of that range: **up to 50%** of HIV-infected individuals can be shown to have cognitive impairment on careful testing. The catch, and it is the clinically decisive one, is that most of these patients are clinically asymptomatic. The impairment may improve on retesting, and patients are often unaware of the changes, which is the opposite of the relentless pre-antiretroviral picture in which cognition tracked steadily downhill.

up to 50%	52%	43%	25-35%
IMPAIRED ON CAREFUL NEUROPSYCHOLOGICAL TESTING	OVERALL IMPAIRMENT, US SIX-CENTRE COHORT	OVERALL PREVALENCE, GERMAN HIV-CLINIC COHORT	IMPAIRMENT IN ART-TREATED COHORTS

GOOD TO REMEMBER

A single positive cognitive screen in a person with HIV is not a diagnosis of dementia. Because up to half test positive on careful assessment yet most are asymptomatic and unaware, and because performance may improve on retesting, a low score is a prompt to grade the deficit and to repeat the assessment, not a verdict. Most of what you detect will be at the asymptomatic or mild end of the spectrum.

2.6 The cohort trap: two distributions, two populations

This is the single most examinable pitfall in the topic, and it is a lesson about reading sources rather than about HIV. Two clean sources give materially different category splits for HAND, and both are correct for the population each describes. A clinically diverse cross-sectional study of 1,555 patients at six academic centers showed an overall rate of neuropsychological impairment of 52%, of which the HAND categories accounted for 47%: 33% ANI, 12% MND and only 2% HAD. A German HIV-clinic cohort gave a different picture: an overall HIV-associated neurocognitive disorder prevalence of 43%, of whom 90% were in treatment, distributed as 20% ANI, 17% mild NCD and 6% HAD.

COHORT AND POPULATION	OVERALL IMPAIRMENT	ANI	MND / MILD NCD	HAD
US six-centre cross-sectional study (n=1,555)	52%	33%	12%	2%
German HIV-clinic cohort (90% on treatment)	43%	20%	17%	6%

TRAP **There is no single canonical HAND distribution to memorise.** The two cohorts agree qualitatively, that milder impairment predominates and frank dementia is now uncommon, but their exact percentages differ by country and case-mix. A stem that asks for "the" proportion in each category is either quoting one specific study or testing whether you will invent an average. Name the population before you quote a figure, and never harmonise the two into a single set of numbers.

2.7 The shift in the treated era

Read together, the figures tell a consistent story about what effective treatment has done to the shape of the disease. In cohorts mostly on effective antiretroviral treatment and assessed with comprehensive batteries, overall rates of cognitive impairment run around 25% to 35%. Within that impaired group the weight has shifted downward: mild neurocognitive disorder accounts for approximately one-quarter, while the criteria for major neurocognitive disorder are met for typically **fewer than 5%** of individuals with HIV-related neurocognitive disturbance. Severe HIV-associated dementia, once the defining neurological complication of advanced disease, is now the exception.

That shift is why the classification matters more now than it did in the past, not less. When almost everyone with cognitive impairment was heading toward dementia, a grading scheme added little; when most impairment is asymptomatic or mild and may not progress, the tier a patient occupies genuinely guides how closely to watch and how hard to look for a reversible contributor. The persistence of impairment despite effective treatment is itself the point of interest: cognition no longer collapses as it once did, but a substantial minority of well-treated patients still test below par, which tells you that suppression of the virus in the blood does not automatically clear the problem in the brain. That residue is what keeps the milder tiers clinically alive and what makes the sub-threshold category worth defining at all. The reasons the treated brain fares better, the biology of the neuro-injury and the role of antiretroviral penetration into the central nervous system, belong to the account of HIV neuro-injury given elsewhere in this chapter and are named here only as the mechanism behind the change.

In the treated era most HIV-related neurocognitive impairment is asymptomatic or mild and frank HIV-associated dementia is uncommon, but the exact category split differs by cohort, so name the population before you quote a figure.

2.8 Assessing and following the patient

The practical corollary of an epidemiology dominated by quiet impairment is that the diagnosis rests on deliberate assessment rather than on spontaneous complaint. Patients are often unaware of the changes, so waiting for a subjective cognitive complaint will miss most cases; screening the at-risk patient, and repeating it, is the only reliable way to catch and stage the deficit. The standardised cognitive instruments used to do this, the bedside and office batteries, are covered in the Psychotherapies, clinical assessment, MSE and nosology notes and are not reproduced here; what this topic adds is the discipline of grading the result against the functional-status hinge rather than reading a single cut-off score as a diagnosis.

Following the patient over time is as important as the first assessment, because the tiers are not fixed. Impairment may improve on retesting, a stable ANI can be an early signal rather than an endpoint, and an intercurrent illness or a new secondary cause can push a patient across a line. The psychiatric consequences of the disease, the depression, and the other comorbidities that so often accompany HIV, are addressed separately in this chapter; the cognitive spectrum described here is one strand of a larger neuropsychiatric picture, not the whole of it. In practice the two strands interact at the bedside, because low mood, poor sleep and anxiety all depress test performance in their own right, and a patient assessed in the middle of a depressive episode can look more impaired than they are. That is one more argument for repeating the assessment rather than fixing a tier on a single sitting, and for treating the reversible before concluding that what remains is HAND.

IN INDIA

Where HIV remains a live clinical entity, keep HAND on the differential for any person living with HIV who reports or shows cognitive change, and treat the diagnosis as one of exclusion: document the infection, then actively rule out the secondary brain diseases before applying the label. Because up to half of patients test positive on careful assessment and most are unaware of it, formal screening of the at-risk patient will pick up impairment that the history alone would miss, and grading it correctly keeps a treatable mimic from being written off as dementia.

3 · The pathophysiology of HIV neuro-injury and the effect of antiretroviral therapy

The single fact that unlocks this whole topic is that HIV does not kill the neurons it damages by living inside them. The brain is injured while the virus sits, for the most part, in other cells altogether. Get that relationship right and the neuropathology, the cognitive profile, the risk markers and the response to treatment all fall into place; get it wrong and every downstream question becomes a memory test you will fail. The marks here reward mechanism rather than lists: how the virus reaches the central nervous system, why the damage it does is indirect, where in the brain that damage concentrates, and what antiretroviral therapy does and does not achieve. The clinical syndromes that this injury produces, the HIV-associated neurocognitive disorders and their formal classification, are set out separately in these notes; here the concern is the pathophysiology that sits underneath them and the effect of treatment on that pathophysiology.

Keep one discipline throughout: the specifics that matter are few, and everything else is reasoning from them.

3.1 How HIV reaches the brain: the Trojan-horse route and the viral reservoir

The virus reaches the central nervous system early, and it does so by hitching a ride inside cells it has already infected. HIV infects several cell types, chiefly the CD4 T-helper lymphocytes and monocytes, and it is the infected monocyte, crossing the blood-brain barrier and maturing into a perivascular macrophage, that carries the virus across. This is the **Trojan-horse** model of central nervous system entry, described in Kaplan and Sadock's Comprehensive Textbook of Psychiatry: infected monocyte-derived macrophages and lymphocytes seed the brain during early infection. Once inside, the virus infects macrophages and microglia, and a small subset of astrocytes may harbour productive infection, as set out in the DSM-5-TR.

What follows is the feature that makes HIV brain disease so hard to eradicate. The virus establishes a locally productive infection that can run relatively independently of viral activity elsewhere in the body, so the brain becomes a **viral reservoir** in its own right. Dulcan's Textbook of Child and Adolescent Psychiatry records how fast this happens: the brain is a repository for HIV, and the virus may penetrate brain tissue as early as **2 weeks** after infection. A compartment that is seeded within a fortnight and thereafter sustains its own replication is a compartment that systemic control alone may never fully reach.

Indirect	CD4 and monocytes	Basal ganglia	Aggressive ART
HOW HIV INJURES NEURONS, NOT BY INFECTING THEM	THE CELLS HIV CHIEFLY INFECTS	HIT EARLY; FRONTOTEMPORAL LOSS COMES LATE	THE BEST AND ESSENTIALLY ONLY TREATMENT FOR HAND

3.2 The central principle: neurons are bystanders, not hosts

The neuron is not the host cell. The effects of HIV on neural functioning are, as Kaplan and Sadock's Comprehensive Textbook of Psychiatry puts it, largely indirect: they arise through exposure to viral proteins, through the proinflammatory cytokines produced by infected or activated monocytes and other cells, and through alteration of the blood-brain barrier, rather than through direct infection of neurons. This is the point on which the whole subject turns, and it is worth stating the comparison explicitly because examiners test exactly this discrimination. The injury is indirect, mediated by proteins, cytokines and a leaking barrier, and it is not the consequence of the virus replicating inside neurons.

The blood-brain barrier belongs in the same account rather than as an afterthought. Its disruption is one of the three routes of indirect injury, and it is also self-amplifying: a leaking barrier admits more infected monocytes, which mature into more perivascular macrophages, which release more of the products that both injure tissue and loosen the barrier further. The same logic extends to the other cells that are damaged. According to Tasman's Psychiatry, neither the oligodendrocytes nor the neurons are themselves infected; both are damaged through indirect routes, by the surrounding macrophages and by the toxicity of viral proteins. This dissociation, between where the virus lives and where the damage shows, is the single organising

idea of HIV neuropathogenesis, and it explains why a patient can carry a heavy brain burden of injury with almost no virus demonstrable in the injured cells at post-mortem.

RULE **HIV injures the brain indirectly; it does not directly infect neurons.** The productively infected cells are macrophages, microglia and a small subset of astrocytes, as recorded in the DSM-5-TR. Neurons and oligodendrocytes are damaged by viral proteins, by macrophage-derived cytokines and by blood-brain-barrier disruption, never by the virus setting up house inside them.

MNEMONIC

Proteins, Cytokines, Barrier

The three arms of indirect neuronal injury in HIV: toxic viral **P**roteins, macrophage-derived proinflammatory **C**ytokines, and disruption of the blood-brain **B**arrier. None of the three requires the neuron itself to be infected.

3.3 The neurotoxic cascade: viral proteins and macrophage-activation products

Unpacking the indirect mechanism is where the mechanistic marks are won. Dulcan's Textbook of Child and Adolescent Psychiatry lists the implicated routes: direct neuronal injury by viral proteins, of which **Tat** and gp120 are the named culprits; the products of macrophage activation, a suite of proinflammatory cytokines and chemokines including interleukin 1, interleukin 6, tumour necrosis factor alpha, interferon gamma, monocyte chemoattractant protein 1, and macrophage inflammatory proteins 1alpha and 1beta; neuroreceptor blockade; autoimmunity; and antibody-mediated cellular toxicity. The through-line is that macrophage and microglial activation, not neuronal infection, drives the chemical injury.

Tasman's Psychiatry sharpens the picture of one protein in particular. The **gp120** glycoprotein, carried on the surface of the virus, is toxic to cells it never infects. Its consequences map onto two compartments of tissue. When oligodendrocytes are lost, a **leukoencephalopathy** emerges, the white-matter injury that gives HIV brain disease much of its imaging signature; and when neurons are touched, they undergo both apoptosis and dendritic simplification, a pruning back of the synaptic tree rather than outright cell death alone. The accompanying figure sets the Trojan-horse entry, the microglial activation and this protein-driven cascade in a single sequence.

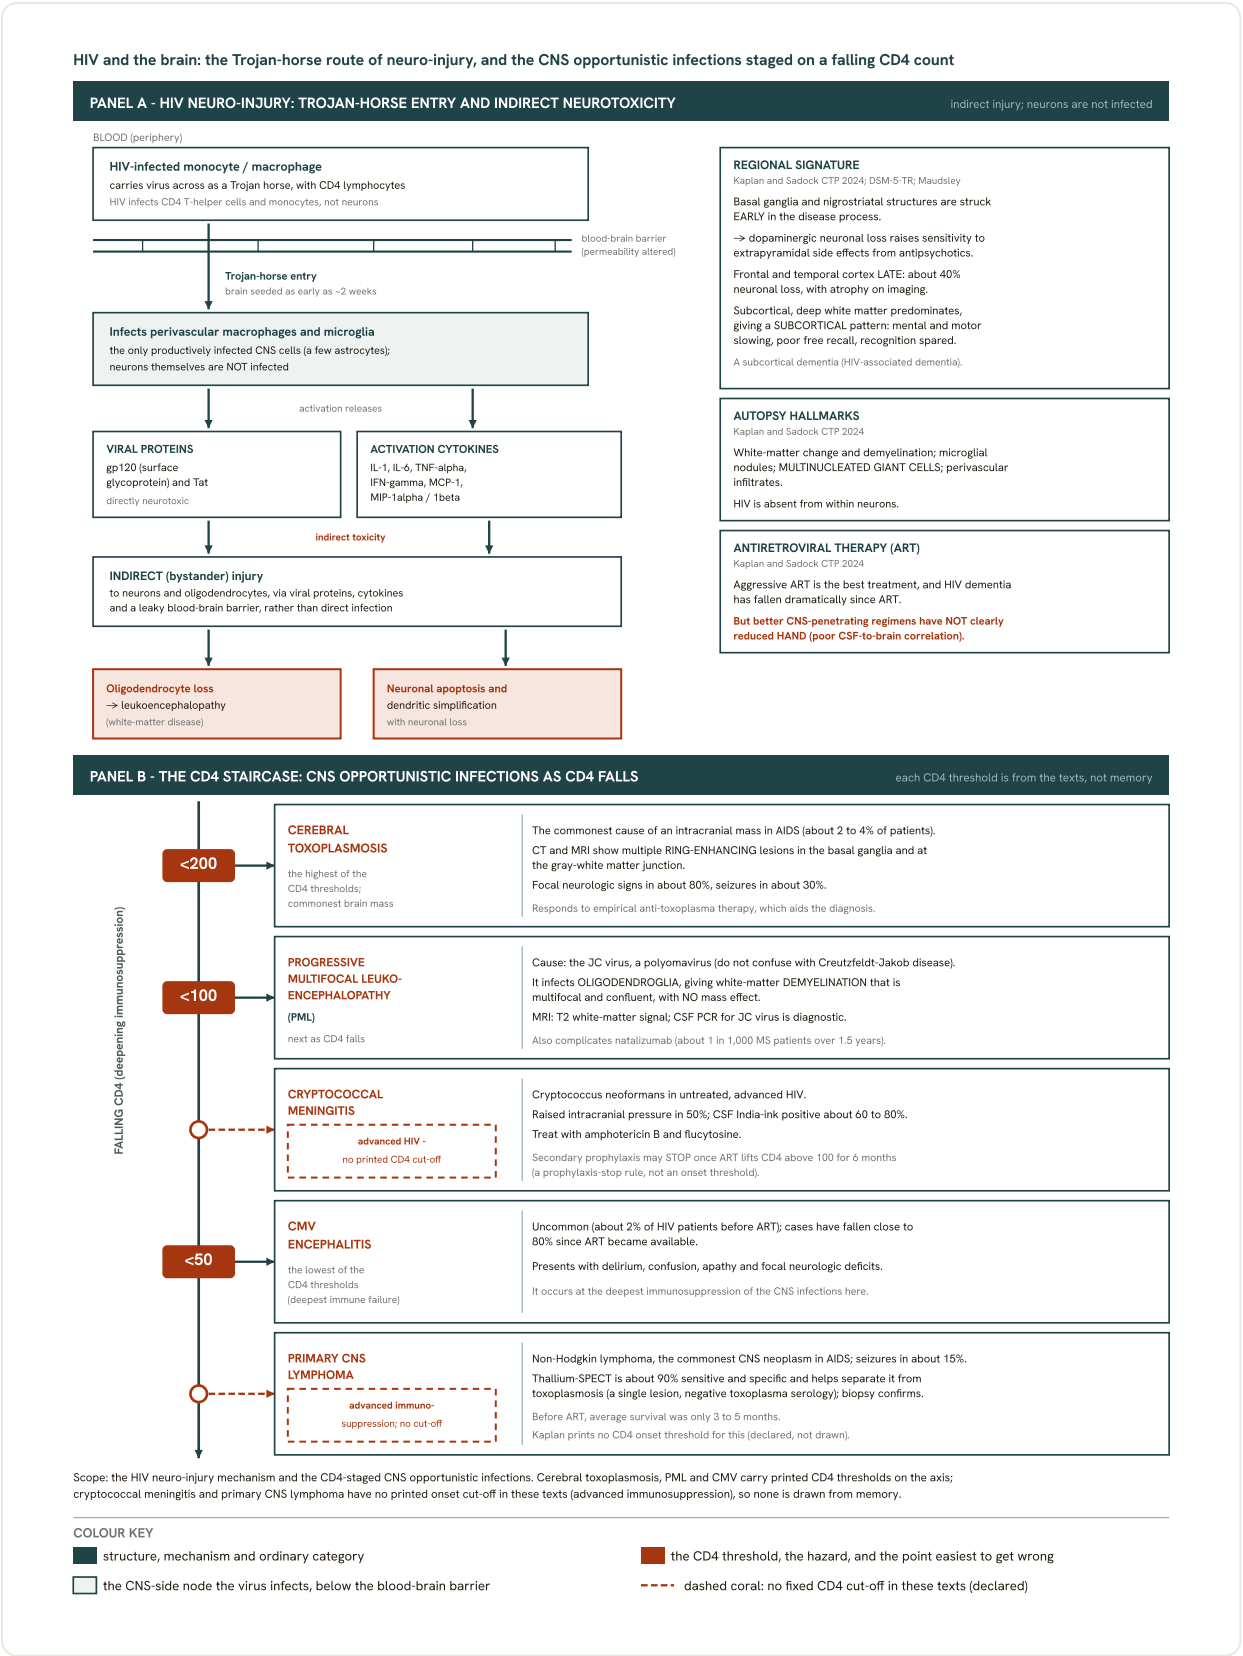


Fig 35.2 HIV neuro-injury and the CD4-staged CNS opportunistic infections. Panel A: the Trojan-horse route into the brain and indirect neurotoxicity that injures neurons without productively infecting them, with the regional and treatment signatures. Panel B stages toxoplasmosis, PML and CMV encephalitis by their printed CD4 thresholds, while cryptococcal meningitis and primary CNS lymphoma carry no printed onset cut-off in these texts. (Kaplan and Sadock CTP 2024; DSM-5-TR; Maudsley 15th edition; Dulcan; Tasman Psychiatry; Kaufman's Clinical Neurology for Psychiatrists.)

CELL TYPE	PRODUCTIVELY INFECTED BY HIV?	ROLE IN THE NEURO-INJURY
Monocyte-derived macrophages	Yes, and the Trojan-horse carrier into the brain	Deliver virus across the barrier; release neurotoxic cytokines
Microglia	Yes	Sustain the central nervous system reservoir; activation drives injury
Astrocytes	A small subset only	Contribute to barrier dysfunction and the inflammatory milieu
Oligodendrocytes	No	Damaged indirectly; their loss yields the leukoencephalopathy
Neurons	No, HIV is markedly absent within neurons	Injured indirectly; apoptosis and dendritic simplification

3.4 The anatomy: a subcortical disease that begins in the basal ganglia

Where the injury concentrates determines what the patient looks like. Kaplan and Sadock's Comprehensive Textbook of Psychiatry describes an early predilection: the basal ganglia and the nigrostriatal structures are affected early in the dementing process, with more diffuse neuronal losses following. The late picture is quantified as an approximate **40%** reduction in frontal and temporal neurons, with atrophy that becomes visible on imaging. The temporal order matters and must not be reversed: deep grey matter first, cortical loss later.

This regional bias produces a recognisable cognitive signature. The DSM-5-TR notes that because HIV preferentially affects subcortical regions over the course of the illness, including the deep white matter, the disorder progresses in a subcortical pattern: mental slowing coupled with motor dysfunction, deficits in procedural learning and in free recall, and relative sparing of recognition memory, verbal abstraction and naming. That profile, slowed and effortful rather than amnestic and aphasic, is the fingerprint of a disease that lives in white matter and deep nuclei. The memory dissociation is worth dwelling on because it is diagnostically useful and frequently examined. A subcortical disorder impairs the effortful, self-generated retrieval that free recall demands, while leaving the stored trace itself relatively intact, so performance recovers markedly when the patient is offered cues or a recognition format. That is the opposite of a cortical amnestic dementia, in which the trace is not laid down and cueing does not help. When the deficit is retrieval rather than encoding, and when motor slowing and procedural difficulty travel alongside it, the anatomy is pointing downward, into the white matter and deep grey, and away from the cortex. The broader distinction between subcortical and cortical dementia is developed in the Dementias and neurocognitive disorders notes, and the standardised instruments used to measure the deficit belong to the Psychotherapies, clinical assessment, MSE and nosology notes; here the point is simply that the anatomy predicts the phenotype.

3.5 The neuropathology at autopsy: a signature you can recognise

The post-mortem findings close the loop between mechanism and syndrome. Autopsy studies of patients who died with AIDS dementia, reported in Kaplan and Sadock's Comprehensive

Textbook of Psychiatry, show a characteristic constellation: white-matter changes and demyelination, **microglial nodules**, **multinucleated giant cells**, and perivascular infiltrates, set against a marked absence of HIV within the neurons themselves. The multinucleated giant cell, formed by the fusion of infected macrophages and microglia, is the histopathological hallmark, and its presence alongside a scarcity of virus in neurons is the tissue-level restatement of the indirect-injury principle.

The same source names the clinical counterpart of these findings. The disorder is a **subcortical dementia**, dominated by loss of attention and concentration and by marked motor slowing. This is the description a candidate should be able to give without hesitation: not a cortical amnestic dementia but a slowed, dysexecutive, motorically retarded picture that follows directly from the deep-structure and white-matter pathology described above.

PITFALL

Reading the psychomotor slowing, poor concentration and flattened output of early HIV brain disease as a depressive episode and stopping there. The subcortical profile of slowed thinking with motor slowing overlaps heavily with the appearance of depression, and the two can coexist; anchoring the assessment to the neurocognitive pattern, and to serostatus, rather than to surface affect keeps a reversible neurological disorder from being managed as a mood disorder alone.

3.6 Dopaminergic vulnerability and its bedside consequence

The early basal-ganglia involvement is not a purely academic point; it changes how you prescribe. The Maudsley Prescribing Guidelines make the clinical consequence explicit: people living with HIV are more susceptible to **extrapyramidal side effects**, because HIV, a neurotropic virus, enters the brain and replicates in the basal ganglia, producing dopaminergic neuronal loss. A brain already depleted of striatal dopamine has less reserve to absorb the dopamine blockade that antipsychotic drugs impose, so movement side effects appear readily and at exposures a seronegative patient would tolerate.

This dovetails with the anatomy established earlier: the same early nigrostriatal injury that begins the dementing process also strips the motor system of its dopaminergic buffer. The practical corollary belongs at the bedside rather than in a mechanism diagram, and it is one of the few places where this pathophysiology directly dictates a prescribing decision.

GOOD TO REMEMBER

When an antipsychotic is genuinely needed in a person living with HIV, expect extrapyramidal effects sooner and at lower exposures than usual, because the virus has already depleted basal-ganglia dopaminergic neurons. Favour an agent with lower extrapyramidal liability, start low and titrate slowly, and watch the motor examination closely. The management of the psychiatric comorbidities themselves, and the interactions between psychotropics and antiretrovirals, are addressed separately.

3.7 What antiretroviral therapy does, and the penetration paradox

Treatment is the part of this story with a genuinely happy chapter and a genuinely humbling one. The happy chapter is that with the arrival of antiretroviral therapy there has been a dramatic reduction in HIV-induced dementia. Kaplan and Sadock's Comprehensive Textbook of Psychiatry is unambiguous about what to do: the best treatment for any HIV-associated neurocognitive disorder is aggressive antiretroviral therapy in patients who are untreated or inadequately treated, and this may produce dramatic improvement. The corollary is sobering in equal measure: to date no clearly effective treatment for these disorders beyond antiretroviral therapy has emerged, so there is no adjunctive neuroprotective agent to reach for once the virus is suppressed.

The humbling chapter concerns drug penetration into the brain, and it is a classic examination trap. The intuitive expectation, that antiretrovirals which penetrate the central nervous system better would be associated with less neurocognitive disorder, has not materialised. The same source offers the reasons: a poor correlation between drug concentrations in the cerebrospinal fluid and those in the brain itself, direct mitochondrial toxicity of the drugs, the degree of immune reactivation, and subtle ongoing viral replication within the central nervous system. Note the direction carefully, because it is the opposite of what feels right: better central penetration has not been shown to buy less injury. Each reason repays a moment's thought. Cerebrospinal-fluid concentration is a convenient surrogate but a treacherous one, since the drug the neuron actually sees is the brain-parenchymal level, which the fluid measures poorly. The drugs themselves are not inert in the compartment they are meant to protect, carrying a mitochondrial toxicity that can injure the very tissue being defended. And a compartment that was seeded within the first fortnight and thereafter replicates on its own terms may sustain low-grade viral activity that suppression of the blood never quite silences. The paradox, then, is not a failure of logic but a lesson in how imperfect a proxy the cerebrospinal fluid is for the brain.

TRAP **Assuming that the more central-nervous-system-penetrant a regimen is, the less HIV neurocognitive disorder it will cause.** This is the expected finding that, per Kaplan and Sadock's Comprehensive Textbook of Psychiatry, has not materialised, plausibly because cerebrospinal-fluid levels track brain levels poorly, the drugs carry their own mitochondrial toxicity, and low-level viral replication grinds on in the compartment. Candidates who have half-remembered a ranking of drugs by brain penetration reach for it here and answer the wrong question.

HIV never directly infects neurons; the brain is injured indirectly, and aggressive antiretroviral therapy in the untreated or inadequately treated patient is the one treatment that can produce dramatic improvement in the neurocognitive disorder.

3.8 Predicting who is at risk: the biomarker and clinical panel

Because the injury is inflammatory and immune-driven, the markers that predict it are markers of immune activation and neuronal breakdown rather than of the virus alone. The DSM-5-TR sets out the panel: the risk of an HIV-related neurocognitive disorder rises with prior episodes of

immunosuppression, with high viral loads in the cerebrospinal fluid, and with a cluster of raised peripheral markers, alongside several clinical and haematological red flags. The single most illuminating of the newer markers is **neurofilament light chain**, a direct readout of axonal injury that ties the blood test back to the tissue damage described throughout this section.

- Immune-activation and inflammatory markers: tumour necrosis factor alpha, interleukin 6, C-reactive protein, D-dimer, sCD14 and sCD163.
- Neuronal-injury marker: neurofilament light chain, rising with axonal damage.
- Virological and immunological history: high cerebrospinal-fluid viral load, a low CD4 cell nadir, and prior episodes of immunosuppression.
- Systemic red flags: anaemia and hypoalbuminaemia.

Read together, these predictors reinforce the mechanistic argument of the whole topic: cumulative immunosuppression and sustained immune activation, not a snapshot of viral load, are what put the brain at risk, which is also why aggressive antiretroviral therapy started before that damage accumulates is the intervention that matters most.

IN INDIA

With a large population living with HIV, the potentially improvable subcortical neurocognitive disorder described here remains a live clinical entity, not a historical one. In any patient presenting with new mental and motor slowing, particularly a younger adult, keep HIV-associated neurocognitive disorder in the differential and establish serostatus early, because aggressive antiretroviral therapy in the untreated or inadequately treated patient can produce dramatic improvement in exactly this picture. The opportunistic infections that complicate advanced disease, and the immune reconstitution that follows treatment, are addressed separately.

4 · Psychiatric comorbidity in HIV: depression, mania, psychosis, suicide and anxiety

A new mood, psychotic or suicidal presentation in a patient living with HIV is rarely a coincidence sitting on top of the infection. The virus, the brain injury it produces, the opportunistic diseases it permits, the drugs used to control it and the psychosocial weight of a chronic transmissible illness all converge on the same clinical picture, and psychiatric comorbidity in HIV is common, consequential and heavily examined for exactly that reason. Depression is the dominant theme; mania and psychosis emerge as the disease reaches the brain; and suicide risk runs above what the measurable risk factors alone would predict. This section works through each in turn: how much more common depression is and how to read its epidemiology without being tripped by the comparator, how the mood and psychotic syndromes track the falling CD4 count, why suicide deserves separate vigilance, and how to treat mania and psychosis in a brain that has become exquisitely sensitive to the very drugs that would help it. The cognitive syndromes of HIV, the CNS opportunistic infections and the interactions between antiretroviral and psychotropic drugs are each addressed separately and are named here only where they bear directly on the psychiatric picture.

4.1 Depression: the commonest comorbidity, and reading its epidemiology

Depression is the single most frequent psychiatric comorbidity in HIV and the one that most reliably degrades the course of the medical illness. Living with HIV carries a three- to fourfold higher prevalence of **major depression** than is seen in the general population. That figure is the one to hold, but it must be read alongside a second, apparently smaller, number that is not in conflict with it. When the comparison is made not against the general population but against an **at-risk HIV-negative comparator**, a group already loaded with the same social adversity, substance use and disadvantage that so often accompany HIV, a meta-analysis of ten studies found a twofold increase in major depression among those infected. The two multipliers answer two different questions. Against the population at large the excess looks larger, because the comparison group carries little of the surrounding hardship; against a well-matched at-risk group the excess shrinks, because much of that hardship is shared, and what remains, the smaller multiplier, is a less-confounded estimate of the association rather than a clean causal measure of the virus and its consequences themselves.

TRAP **Reading the two- and the three- to fourfold depression figures as a contradiction.** They are not in conflict. The larger multiplier is measured against the general population and the smaller against an at-risk HIV-negative comparator, so the smaller figure is not a weaker version of the larger one but a cleaner estimate of the disease-attributable excess once shared social adversity has been controlled for. A stem that quotes one and asks you to reconcile it with the other is testing whether you noticed that the denominator changed.

3- to 4-fold	2-fold	2.5-fold
MORE DEPRESSION THAN THE GENERAL POPULATION	VERSUS AN AT-RISK HIV-NEGATIVE COMPARATOR	AS CD4 FALLS, NEAR THE AIDS TRANSITION

4.2 Depression tracks the failing immune system

Depression in HIV is not distributed evenly across the course of the illness; it intensifies as immune function collapses. The **Multicenter AIDS Cohort Study** demonstrated a two-and-a-half-fold rise in rates of depression as the CD4 count fell **below 200**, in the window just before patients crossed into AIDS. That gradient is clinically instructive, and the psychiatric burden shifts as immune function declines. A depressive syndrome appearing for the first time in a patient whose counts are dropping is not simply a psychological reaction to bad news. It coincides with advancing immunosuppression and a mounting systemic and central-nervous-system disease burden, the virus having seeded the brain far earlier in the infection, and it should draw as much attention to the state of the infection as to the mood itself.

The gradient also reshapes the differential. Early in the illness, reactive and social factors dominate, and a low mood is most readily understood as an adjustment to diagnosis, stigma and uncertainty. Late in the illness, once the count has fallen and the brain is involved, the same syndrome carries a heavier biological loading and sits closer to the neuropsychiatric complications of advanced disease. A first depressive episode at that stage therefore warrants a wider look than the same complaint would at seroconversion, because the machinery generating it has changed even when the surface presentation has not.

4.3 Why the depression matters: adherence, transmission and outcome

Depression earns its place at the front of this topic because it alters the trajectory of the medical illness, not merely the patient's comfort. It exerts a negative influence on medical care and outcome, and it does so through two channels in particular. Patients who are depressed take their antiretroviral therapy less reliably, so virological control slips; and depression is associated with more **transmission-risk behaviour**, raising the chance of onward spread. The relationship runs in both directions, because advancing infection and its brain involvement in turn deepen the mood disturbance, so that a treatable psychiatric problem and a progressing medical one feed each other if neither is addressed.

Anxiety travels alongside the depression in this population and shapes engagement, sleep and the experience of somatic symptoms, though the strength of that association is not quantified in the source used here and is not asserted as a figure. What can be said with confidence is that recognising and treating the affective burden is inseparable from managing the infection well. The primary nosology and pharmacotherapy of the mood disorder itself belong to the Mood disorders: pharmacotherapy notes and are not reproduced here; what matters in the HIV clinic is that the mood syndrome is actively looked for rather than waited for.

> **GOOD TO REMEMBER**
>
> Screen for depression at every review of a patient on antiretroviral therapy, and treat it as a driver of the medical result rather than as a separate complaint. Because depression worsens adherence and outcome, recognising and treating it is part of protecting the virological response, not an optional extra. In practice the depressed, poorly adherent patient and the patient with rising viral load are frequently the same person.

4.4 AIDS mania: how often, and the two populations it comes from

Mania is far less common than depression but far more striking, and it carries a distinctive relationship to disease stage. In one clinic series, **AIDS mania** was recorded in **8%** of all patients with AIDS seen across seventeen months of attendance. Set against the baseline, that is **more than ten times** the six-month prevalence of mania in the general population, which is why the syndrome is treated as a genuine complication of the disease rather than a chance co-occurrence.

Two populations sat within that group, separated by their CD4 count and by the timing of their first manic episode. One set had its first episode early, while CD4 counts were still above 200; the other had its first episode late, once CD4 counts had fallen below 200. The distinction is not bookkeeping. The early group behaves like people with a pre-existing bipolar diathesis whose illness happens to surface during the infection, whereas the late group is where the mania is being generated by the disease in the brain. The two look superficially alike at the bedside and pull apart sharply once history, family history and cognition are examined, which is the subject of the next section.

4.5 The face of HIV-driven mania

The late-onset group is the one that matters most to the neuropsychiatrist, because its mania behaves like a **secondary mania**, a mood syndrome arising from brain disease rather than from

an inborn mood diathesis. It is largely confined to late HIV infection and appears to be a consequence of advanced HIV-related brain involvement. These patients are less likely to carry a personal or family history of mania or mood disorder, and more likely to show dementia or other cognitive impairment alongside the mood change; the co-occurring cognitive syndromes of HIV are addressed separately and are named here only as company the mania keeps. The mood itself tends to be irritable rather than euphoric, the course chronic rather than episodic, and the illness often severe and malignant in its trajectory, which is a different and more forbidding picture than the classic euphoric mania of primary bipolar disorder.

The clinical importance of naming it as secondary is twofold. First, it reorients the assessment: rather than reaching for a mood-disorder formulation and a longitudinal history of past episodes, the clinician who recognises the pattern goes looking for the state of the brain and the immune system that produced it. Second, it warns about the course. A chronic, irritable, cognitively coloured mania that does not remit in the way an episodic bipolar illness would is both harder to live with and harder to treat, and the malignant trajectory means that undertreatment and overtreatment are both dangerous. The bedside discriminators that matter most are the absence of a prior mood history, the presence of cognitive change, and the timing of onset deep into the infection; taken together they point away from a primary diagnosis and towards the diseased brain.

FEATURE	EARLY-ONSET MANIA	LATE-ONSET, HIV-DRIVEN MANIA
Timing and CD4	Early, CD4 above 200	Late, CD4 below 200
Personal or family history of mood disorder	More likely than in late-onset; suggests a primary bipolar diathesis	Characteristically absent
Cognitive state	Usually less impaired than in late-onset	Dementia or cognitive impairment common
Predominant mood quality	More often euphoric or classic	Irritable more than euphoric
Course	More often episodic	Chronic, often severe and malignant

MNEMONIC

Irritable, Chronic, No family history, Cognitive decline

The signature of HIV-driven late-onset mania: an **I**rritable rather than euphoric mood, a **C**hronic rather than episodic course, **N**o personal or family history of mood disorder, and accompanying **C**ognitive decline, the whole picture appearing only once the infection is advanced.

4.6 New-onset psychosis and mania as late-stage events

Psychosis in HIV, when it appears for the first time in advanced disease, should be read the same way as the late mania: as a signal from the brain rather than the start of an independent psychiatric career. **HIV-associated psychosis** and, as a distinct syndrome, new-onset mania can each emerge in late-stage disease, arising as part of the HIV disease process itself or from

opportunistic infections in the central nervous system, and both may respond to low-dose neuroleptics. The practical consequence is a single reflex: a first psychotic or manic episode in this setting is a secondary syndrome until the workup says otherwise, and the search is for a neurological cause before a psychiatric label.

RULE **A first manic or psychotic episode in advanced HIV is secondary until proven otherwise.** New-onset mania or psychosis late in the illness can be a manifestation of HIV disease or of a CNS opportunistic infection, so it demands a search for a neurological cause before it is treated as a primary psychiatric disorder. The primary nosology and management of psychotic illness belong to the Schizophrenia: management, antipsychotics and adverse effects notes.

Holding that rule means keeping a short differential in view whenever the picture is new:

- The HIV disease process reaching the brain in advanced infection.
- CNS opportunistic infections, which are addressed separately and not detailed here.
- An intercurrent delirium from systemic illness, medication or metabolic disturbance.
- An underlying primary mood or psychotic disorder surfacing more or less independently of the infection.

4.7 Suicide: a risk above the sum of its parts

Suicide in HIV is not fully explained by the depression, the substance use and the social adversity that accompany it. Suicidal ideas, suicidal behaviour and **completed suicide** are all increased in people living with HIV, and the crucial observation is that this excess appears to persist beyond the contribution of the known mediating risk factors. Even after accounting for the measurable drivers, in other words, HIV seems to add risk of its own. That is the point to carry into any assessment: a patient whose depression is mild and whose social situation looks manageable is not thereby low-risk simply because the familiar predictors are quiet.

Treatment has changed the landscape without flattening it. The development and use of antiretroviral therapy has been associated with a fall in suicide risk, a welcome shift that tracks the broader improvement in prognosis, but the risk has not returned to that of the background population. The residual elevation is the part that matters clinically, because it is easy to assume that a patient doing well on treatment has been returned to baseline on this axis too, when the evidence says otherwise.

PITFALL

Assuming that a virologically well-controlled patient on antiretroviral therapy no longer carries an elevated suicide risk. Treatment lowers the risk but does not normalise it, so good viral control is not a reason to relax suicide assessment. The stable, adherent, outwardly-thriving patient still warrants direct and unhurried enquiry about suicidal thinking.

4.8 Treating the mania and psychosis safely

Treatment of the mood and psychotic syndromes of advanced HIV is governed by one overriding fact: the same brain that has turned manic or psychotic has also become intolerant of the drugs

used to treat it. Most care happens in the outpatient setting shared with HIV medicine, with admission reserved for the patient whose mania or psychosis threatens safety or is being driven by an acute intercurrent illness, and the emergency presentation is the acutely disturbed, disinhibited patient who cannot be contained at home. The immediate target is behavioural control and the treatment of any precipitating infection or metabolic disturbance; over the medium term the aim is remission of the syndrome together with protection of cognition and of adherence to the antiretroviral regimen. Across every mode of treatment, drug, psychological and social, the governing instruction is caution with dose. The psychological work is largely supportive and practical, aimed at containment, orientation and the rebuilding of adherence, rather than the insight-oriented approaches that a cognitively impaired and acutely disturbed patient cannot use. The social mode carries real weight here, because housing, income, stigma and the reliability of a caregiver determine whether the antiretroviral regimen survives the episode at all.

Pharmacologically, patients with AIDS mania respond to relatively low doses of antipsychotic medication and are very sensitive to their side effects, even with second-generation agents; new-onset psychotic symptoms similarly respond to low-dose neuroleptics. The sensitivity does not stop at the antipsychotic itself. It extends to the drugs used to manage adverse effects, so that extrapyramidal symptoms treated with even modest doses of anticholinergic agents can precipitate an **anticholinergic delirium** in a brain that is already cognitively vulnerable. The interactions between antiretroviral and psychotropic drugs, which further constrain which agent and which dose are safe, are addressed separately and must be checked before prescribing.

IN INDIA

Where HIV remains a live part of the differential, a first manic or psychotic episode presenting in mid-adult life with cognitive slowing, or with an atypical, irritable and chronic course, should prompt review of serostatus and immune function before it is settled as a primary mood or psychotic disorder. Keeping HIV-driven secondary mania and psychosis on the list changes the workup and the prescribing, because both the treatment approach and the prognosis differ from those of a primary illness.

In advanced HIV, treat new mania and psychosis as secondary syndromes of a diseased brain: use low doses of antipsychotic medication, expect heightened sensitivity to side effects even with second-generation agents, and be wary of anticholinergics, which can tip an already vulnerable patient into delirium.

5 · CNS opportunistic infections and HIV-related secondary psychiatric syndromes

In advanced HIV the brain becomes vulnerable to a defined, ordered set of opportunistic infections and tumours, and the order is dictated by the CD4 count. As cell-mediated immunity fails, each fall in the CD4 count opens a fresh window for a different organism, so the count is the single most useful piece of information at the bedside: it tells you which diseases are plausible before the first image is obtained. These processes matter to the psychiatrist for two

reasons. The first is that several of them present with exactly the symptoms a psychiatrist is asked to assess, apathy, low mood, confusion and personality change, so an opportunistic CNS process is a standing differential for what looks like depression or delirium in an HIV-positive patient. The second is that getting the diagnosis wrong is costly, because the treatments diverge sharply and some of these conditions are rapidly fatal if missed. The marks concentrate here because the topic rewards a clean mental model: pair each condition with its CD4 threshold, its imaging signature and its confirmatory test, and hold firmly to the one differential that trips candidates most, the enhancing mass lesion. This account works through the mass lesions first, then the white-matter and meningeal infections, and closes on the secondary psychiatric syndromes these conditions generate. The pathophysiology of HIV neuro-injury itself, and the direct CNS effects of antiretroviral therapy, are taken up separately in these notes and are only named here.

5.1 The CD4 count as the organising axis

The clinical power of the CD4 count is that it stratifies risk into a staircase. Antiretroviral therapy has flattened that staircase, because restoring the CD4 count above the danger thresholds withdraws the immunological permission these organisms need, and the incidence of every disease in this section has fallen where treatment is available. What remains true, and what the examiners test, is the staging logic: a patient who is virologically controlled with a reconstituted immune system is at low risk, whereas the patient who presents late, defaults from treatment, or does not yet know their status arrives with the profound immunosuppression in which these diseases flourish. The strip below sets out the ladder from the least to the most profound immune failure.

< 200	< 100	< 50	up to 10%
CD4 CELLS/MICROLITRE: CEREBRAL TOXOPLASMOSIS BECOMES LIKELY	CD4 CELLS/MICROLITRE: PML TYPICALLY APPEARS	CD4 CELLS/MICROLITRE: CMV ENCEPHALITIS, THE DEEPEST SUPPRESSION	OF AIDS PATIENTS: PRIMARY CNS LYMPHOMA, THE COMMONEST CNS NEOPLASM

Two of these conditions resist a single fixed number and are best held as diseases of advanced disease rather than of a named threshold. Cryptococcal meningitis is described as occurring in untreated patients with advanced HIV, and primary CNS lymphoma is likewise a disease of profoundly low counts for which a defining onset cut-off is not securely established in this source. Writing a precise number for either from memory is exactly the kind of confidently-held error this field punishes, so they are staged qualitatively.

CONDITION	IMAGING OR LESION	CONFIRMATORY POINTER	DISCRIMINATING FEATURE
Cerebral toxoplasmosis	Multiple ring-enhancing lesions in the basal ganglia and at the gray-white junction	Toxoplasma serology; response to empirical anti-toxoplasma treatment	Commonest intracranial mass; multiple lesions, often febrile
Primary CNS lymphoma	Often a solitary enhancing lesion	Thallium SPECT; brain biopsy for confirmation	Afebrile, toxoplasma IgG negative, single lesion
PML	Multifocal white-matter demyelination without mass effect	CSF JC virus PCR	Subacute, progressive, no fever, no enhancement
CMV encephalitis	Clinical: delirium, confusion, apathy and focal deficits	Marker of the deepest immune failure	The lowest CD4 setting of the group
Cryptococcal meningitis	Meningitis with raised intracranial pressure	India ink stain; fungal culture	Subacute meningism; raised pressure the immediate danger

RULE **In a patient with known or suspected advanced HIV, a new neurological or neuropsychiatric change is an opportunistic CNS process until it is excluded.** Stage by CD4 count, neuroimage, and where safe examine the CSF, before a psychiatric label is applied. Reaching for the psychiatric diagnosis first is the single error the whole approach is built to prevent.

5.2 Cerebral toxoplasmosis: the commonest mass lesion

Cerebral toxoplasmosis is reactivation of latent *Toxoplasma gondii* within the brain, and in AIDS it is the single commonest reason for an intracranial mass. Active disease generally occurs once the patient is immunocompromised, with a CD4 count below **200 cells/microlitre**, and across the AIDS population it affects between 2% and 4% of patients. Because the organism seeds the brain through the bloodstream it tends to lodge at vascular watersheds, which is why the lesions cluster in the basal ganglia and at the gray-white matter junction and why they are characteristically multiple.

On CT and MRI the hallmark is multiple **ring-enhancing lesions**: a necrotic centre, a rim that takes up contrast, and surrounding vasogenic oedema that produces mass effect. The ring reflects the immune response mounting a wall around the infection, so the appearance is essentially that of a group of abscesses. Clinically the picture is that of one or more expanding brain lesions. Focal neurologic signs are present in approximately 80% of cases, reflecting the location of the deposits, and focal or generalized seizures occur in roughly 30%. Headache, fever and a subacute encephalopathy commonly accompany them. The reason toxoplasmosis heads the differential for any enhancing lesion in AIDS is simply arithmetic: it is both the commonest such lesion and the most treatable, so it is the diagnosis to assume and treat while the alternatives are being excluded.

5.3 Primary CNS lymphoma and the enhancing-lesion trap

Primary CNS lymphoma, which is a non-Hodgkin lymphoma of the brain, is the commonest CNS neoplasm seen in AIDS, affecting up to 10% of patients, and seizures are present in about 15%. It is the great mimic of toxoplasmosis, because it too presents as an enhancing mass in an immunosuppressed brain, and distinguishing the two is where candidates most often come unstuck. The imaging and clinical clues run in a consistent direction. Lymphoma tends to produce a single lesion in a patient who is afebrile and whose toxoplasma IgG is negative, whereas toxoplasmosis is typically multiple and febrile with positive serology. **Thallium SPECT** is around **90% sensitive and specific** for lymphoma and helps separate the two, but definitive diagnosis rests on brain biopsy. The stakes are high: before antiretroviral therapy the disease carried an average survival of only 3 to 5 months.

TRAP **Treating every enhancing lesion in AIDS as toxoplasmosis and forgetting the lymphoma hiding behind it.** The reflex to cover toxoplasmosis first is correct, because it is commoner and treatable, but the lesion that is solitary, afebrile and toxoplasma IgG negative is the one that is more likely lymphoma. A lesion that fails to respond to an adequate empirical anti-toxoplasma trial should prompt thallium SPECT and biopsy, not a longer course of the same treatment.

GOOD TO REMEMBER

The clean discriminators to carry into the exam and the ward are number, fever and serology. Multiple lesions, fever and a positive toxoplasma titre point to toxoplasmosis; a single lesion in an afebrile, toxoplasma IgG negative patient points to lymphoma and to the need for biopsy. Note that the source used here does not attach a CD4 threshold or a specific oncogenic virus to primary CNS lymphoma, so neither is asserted.

5.4 Progressive multifocal leukoencephalopathy and the JC virus

Progressive multifocal leukoencephalopathy is a demyelinating disease of the white matter that occurs in immunocompromised patients, and in AIDS it typically affects those with fewer than **100 CD4 cells/microlitre**. Its prevalence before antiretroviral therapy was between 1% and 10% of AIDS patients. The causative agent is a polyomavirus, the **JC virus**, named after the patient in whom it was first identified and not, despite the shared initials, to be confused with Creutzfeldt-Jakob disease. That single point of nomenclature is a favourite examination pitfall, and getting it wrong signals a shaky grasp of the whole topic.

The mechanism is worth holding precisely because it explains the imaging. The JC virus targets **oligodendroglia**, the cells that generate and maintain CNS myelin, and through their destruction the infection produces demyelination in the white matter of the brain and spinal cord, along with death of astrocytes. The result is multiple, widespread, often confluent white-matter lesions with a multifocal distribution. On MRI these show as areas of high signal (hyperintensity) on T2-weighted and FLAIR images, with corresponding low signal on T1, primarily in the white matter, and crucially they carry no mass effect, which is what separates PML from the enhancing, space-occupying lesions of toxoplasmosis and lymphoma. Diagnosis is secured by detecting JC virus DNA in the CSF by PCR, a test that is both sensitive and specific. Clinically the disease advances

over weeks as a progressive accumulation of focal deficits, cognitive decline and, not infrequently, the behavioural and affective change that brings it to psychiatric attention.

5.5 PML beyond HIV: natalizumab and reactivation

PML is not exclusive to HIV, and understanding why it appears in other settings sharpens the concept that it is fundamentally a disease of failed immune surveillance over a latent virus. It complicates a range of immunosuppressive situations, including the treatment of multiple sclerosis and organ transplantation. The most instructive example is the monoclonal antibody natalizumab, used in multiple sclerosis: probably as a result of the immunosuppression it produces, roughly 1 person for every 1000 patients treated for 1.5 years initially developed PML.

What subsequent study of these cases revealed is the key teaching point. The multiple sclerosis patients who went on to develop PML had anti-JC virus antibodies present before natalizumab was started, which indicates that the virus was already there, quiescent, and that the drug permitted reactivation rather than fresh infection. The same principle underlies PML in AIDS: the falling CD4 count lifts the immune restraint that had kept a latent polyomavirus in check. Multiple sclerosis and its disease-modifying therapy are addressed in their own right elsewhere in these notes; here natalizumab is named only to make the reactivation mechanism explicit.

5.6 Cytomegalovirus encephalitis: the deepest immunosuppression

Cytomegalovirus encephalitis sits at the floor of the immunological staircase. CMV infection is an uncommon complication of AIDS overall, affecting about 2% of HIV patients before antiretroviral therapy, and clinically evident encephalitis is rarer still, occurring almost only in patients whose CD4 count has fallen below 50 cells/microlitre. It is therefore a marker of the most profound immune failure in this section, and its presence tells you the patient's defences are essentially exhausted.

The clinical syndrome is one that lands squarely in psychiatric territory. The commonest features are delirium, confusion, apathy and focal neurologic deficits, so a CMV encephalitis can present as an acute confusional state or as a flat, withdrawn patient who is easily mistaken for someone who is depressed. Antiretroviral therapy has had a striking effect on its epidemiology, with close to an 80% reduction in cases since effective treatment became available, so like the rest of this group it is now chiefly a disease of the untreated and the late-presenting. When it does appear, it signals that the immune system has reached the point at which the full range of opportunistic processes must be actively sought.

5.7 Cryptococcal meningitis: raised pressure and the India ink stain

Cryptococcal meningitis, caused by *Cryptococcus neoformans*, is the classic fungal meningitis of advanced HIV and is usually seen in untreated patients with advanced disease. Its incidence is now around five cases per 1,000 people living with HIV. It presents subacutely, with headache, fever and meningism that can be surprisingly muted, and the immediate danger is mechanical rather than infective: intracranial pressure is elevated in 50% of patients, and it is the raised pressure, causing visual loss and herniation, that kills early. Recognising and relieving it, by therapeutic lumbar puncture, is as important as the antifungal treatment itself.

The CSF is diagnostically rich but can mislead, because it is normal in about 20% of cases. When abnormal it typically shows a monocytosis, an elevated protein, a decreased glucose and positive fungal cultures. The bedside test worth memorising is the **India ink stain** of the CSF, which is positive between **60% and 80%** of the time and reveals the encapsulated yeast directly. Treatment requires amphotericin B and flucytosine. Once the patient has received antiretroviral therapy for six months and the CD4 count has risen above 100 cells/microlitre, secondary prophylaxis may be terminated: a figure that belongs to stopping prophylaxis after immune recovery, and not to the count at which the disease begins.

- Elevated opening pressure, raised in half of patients and the immediate threat to sight and life.
- A monocytic (lymphocytic) pleocytosis rather than the neutrophil predominance of bacterial meningitis.
- Elevated CSF protein.
- Decreased CSF glucose.
- Positive fungal cultures, with a positive India ink stain in the majority.

IN INDIA

Where late-presenting HIV remains common, cryptococcal meningitis is a leading cause of subacute meningitis in the immunosuppressed adult, so keep it firmly in the differential for any HIV-positive patient with headache and altered behaviour. Two clinical decisions follow directly from the biology above: measure and actively manage the opening pressure at lumbar puncture, because raised pressure is the early killer, and send CSF for an India ink stain and fungal culture rather than treating the presentation as a primary psychiatric disturbance.

5.8 When the infection wears a psychiatric mask

The reason all of this belongs in a psychiatry chapter is that these conditions imitate the disorders a psychiatrist is trained to treat. Several of the HIV-related CNS conditions can produce depressive symptoms in their own right, including toxoplasmosis, cryptococcal meningitis, lymphoma and syphilis. More insidiously, AIDS dementia and other HIV-related CNS conditions can generate a flat, apathetic state that is often misdiagnosed as depression, and a superimposed delirium can mimic a wide range of psychiatric presentations. The clinical consequence is that low mood, apathy or confusion in an HIV-positive patient is never to be taken at face value as a primary mood disorder until an organic cause has been considered and, where indicated, excluded. The primary nosology and pharmacotherapy of depression itself belong to the Mood disorders: pharmacotherapy notes, and the assessment and phenomenology of delirium to the Stroke, TBI, delirium and focal brain syndromes notes; the point here is only that this population reaches those diagnoses through a longer and more suspicious route.

One infection in this group deserves special mention as a psychiatric imitator. CNS syphilis, a condition that had become quite rare, has been reappearing in medical centres with HIV specialty services, and it remains **the great imitator** it was originally named, capable of presenting with

depression among a broad range of neuropsychiatric pictures. Its full staging, features and penicillin treatment are set out in the dedicated account of neurosyphilis elsewhere in these notes; the reason to name it here is that a resurgent syphilis in an HIV clinic is precisely the sort of treatable cause a psychiatrist must not overlook when the presenting complaint is a change in mood or personality.

MNEMONIC

Toxo, Crypto, Lymphoma, Syphilis

The HIV-related CNS conditions that can masquerade as, or produce, depressive symptoms: **T**oxoplasmosis, **C**ryptococcal meningitis, **L**ymphoma and **S**yphilis, with AIDS dementia adding the flat, apathetic state that most often gets misread. Each is a treatable diagnosis hiding behind a psychiatric one.

PITFALL

Labelling the flat, withdrawn, poorly motivated HIV patient as depressed and starting an antidepressant, when the apathy is the face of an AIDS dementia or another CNS condition misdiagnosed as depression. The error is understandable, because the surface presentations are identical, but it delays the neurological workup that the low CD4 count demands and lets a treatable infection or a dementing process advance unrecognised.

In advanced HIV, treat any new focal deficit, seizure, or apathetic-depressive change as an opportunistic CNS process until CD4 staging, neuroimaging and CSF have excluded one; the enhancing mass is toxoplasmosis until biopsy proves lymphoma, and the flat, apathetic patient may harbour a CNS infection rather than a mood disorder.

6 · Antiretroviral-psychotropic interactions and the neuropsychiatric effects of antiretrovirals

A patient living with HIV who also needs a psychotropic is being asked to take two of the most interaction-prone drug groups in medicine at the same time. That single fact organises the whole topic. The interface between antiretroviral therapy and psychiatric medication generates trouble in two directions at once. In one direction the two drug classes interfere with each other's handling and effects, so that a regimen that is safe alone becomes toxic or ineffective in combination. In the other direction several antiretrovirals are themselves psychoactive, producing depression, disturbed sleep, psychosis and suicidal thinking regardless of any co-prescribed drug. The marks here are safety marks, because getting them wrong harms a patient who came for help with the very symptoms your prescribing may cause or worsen.

The depression, mania, psychosis, suicidality and anxiety that complicate HIV are substantial in their own right and are addressed separately in this chapter. The concern here is narrower and more practical. It is, first, the pharmacology of co-prescribing antiretrovirals with psychotropics,

and second, the direct neuropsychiatric toxicity of the antiretroviral agents themselves, of which efavirenz is the exemplar the examiners return to.

6.1 Why the antiretroviral-psychotropic interface carries marks

The starting principle is that **pharmacokinetic interactions** between antiretroviral drugs and psychotropics occur frequently and can be clinically significant. They are numerous and complex, too many to hold in memory reliably, which is why prescribers are directed to regularly updated resources such as www.hiv-druginteractions.org rather than to a fixed list committed to recall. Alongside these, **pharmacodynamic interactions** may also occur through overlapping adverse effects, where two drugs that each nudge the same physiological parameter combine to push it into dangerous territory even when neither has altered the other's blood level.

Holding those two mechanisms apart is the organising move of the topic. A pharmacokinetic interaction changes how much drug is present: one agent alters the absorption, metabolism or clearance of another, most often through hepatic enzymes, so the measured concentration rises or falls. A pharmacodynamic interaction leaves the concentrations untouched and instead sums the effects of two drugs acting on the same target system, the heart's repolarisation or the seizure threshold being the classic examples. The reason the distinction earns marks is that the two demand different safeguards. A pharmacokinetic problem is anticipated by knowing the metabolic route and, where needed, adjusting the dose or checking a level; a pharmacodynamic problem is anticipated by knowing what each drug does to the end organ and monitoring that organ directly.

2-4 weeks	30% lower	2-fold higher
UNTIL THE ACUTE EFAVIRENZ NEUROPSYCHIATRIC SYMPTOMS USUALLY SUBSIDE	EFFECT OF THE INDUCER EFAVIRENZ ON A SHARED SUBSTRATE (LEDIPASVIR, A NON-PSYCHOTROPIC EXEMPLAR)	EFFECT OF THE ENHANCER COBICISTAT ON THE SAME SUBSTRATE

6.2 The cytochrome P450 axis: boosters raise, the inducer lowers

Most of the important pharmacokinetic interactions run through hepatic cytochrome P450, and the CYP3A subfamily in particular. Two antiretroviral strategies sit at opposite poles of this enzyme system, and the exam wants you to know which way each one pushes. **Ritonavir** and cobicistat are boosting agents; cobicistat is best understood as a pure **pharmacokinetic enhancer**, a drug added not for antiviral effect but to inhibit CYP3A and thereby raise the level of the agent it accompanies. By that same inhibition they raise the levels of co-administered psychotropics that are CYP3A substrates. At the other pole, efavirenz behaves as a CYP3A **enzyme inducer**, upregulating the enzyme and lowering the concentration of its substrates.

The magnitude and direction are worth anchoring with a concrete, if deliberately non-psychiatric, illustration. Using the hepatitis C drug ledipasvir purely as an exemplar substrate, the inducer efavirenz reduces its concentration by around **30%**, whereas the enhancer cobicistat increases the same drug's concentration roughly **2-fold**. Ledipasvir is not a psychotropic and these figures do not transfer to any particular psychiatric drug; they are quoted only to fix the two directions in the mind. The clinical corollary is what matters: a boosted regimen and an

efavirenz-based regimen pull a shared substrate in opposite directions, so the same psychotropic may need a higher maintenance dose alongside one and a lower one alongside the other.

RULE **Boosters raise the substrate, the inducer lowers it.** The pharmacokinetic enhancers ritonavir and cobicistat inhibit CYP3A and raise the levels of psychotropic substrates, while efavirenz induces CYP3A and lowers them. Because the interactions are individually numerous and complex, do not trust memory: confirm the direction and magnitude against a maintained interactions resource before you co-prescribe.

6.3 A worked example: when inhibition raises a substrate to toxicity

The danger of the boosting strategy is easiest to see in a single worked interaction. Ritonavir can decrease the metabolism of ecstasy (MDMA) and precipitate toxicity: by inhibiting the enzyme that would ordinarily clear the recreational drug, the antiretroviral lets it accumulate to harmful levels. The principle generalises well beyond one recreational substance. Any cytochrome P450 substrate a patient is taking, prescribed or not, becomes a candidate for the same accumulation once a potent inhibitor is added, and the psychotropics most vulnerable are those with a narrow gap between a therapeutic and a toxic concentration.

What makes this counter-intuitive is that the enzyme inhibition is not an accident of the drug but its intended action. Ritonavir earns its place in modern regimens precisely because it inhibits metabolism and so sustains the levels of the antiretrovirals it accompanies; the same inhibition that boosts a partner drug on purpose boosts an unrelated co-substrate by the same route. The pharmacology working as designed for the antiviral is the pharmacology working against the psychiatrist.

Two habits follow from this. The first is to take a full drug history, including recreational and over-the-counter agents, because the inhibitor is silent until the co-substrate reaches toxicity and gives no warning at the point of prescribing. The second is to treat the boosted regimens with particular respect: they are the regimens most likely to turn a previously well-tolerated psychotropic dose into an overdose without a single change to that psychotropic's prescription. Anticipation, not reaction, is the whole game.

TRAP **Assuming that all antiretroviral therapy lowers psychotropic levels.** Candidates who have learned that efavirenz is an inducer generalise the reflex and expect every regimen to reduce a co-prescribed drug, so they instinctively push the psychotropic dose up. The boosted regimens do the opposite: ritonavir and cobicistat raise the substrate, and the same upward dose adjustment then risks the accumulation that lets ritonavir precipitate frank toxicity. Direction depends on the regimen, and the two are exactly opposite.

6.4 Pharmacodynamic overlap: QT prolongation and the seizure threshold

Not every dangerous interaction changes a blood level, and two pharmacodynamic overlaps are worth committing to memory. The first is cardiac. Several antiretrovirals prolong the **QT interval** in their own right, efavirenz, rilpivirine and the saquinavir-ritonavir combination among them, while atazanavir, lopinavir and saquinavir prolong the PR interval instead. Add a QT-prolonging psychotropic to a QT-prolonging antiretroviral and the two contributions to

repolarisation summate, so the arrhythmia risk compounds. This is precisely why ECG monitoring is recommended when citalopram or escitalopram is co-administered with antiretroviral therapy that prolongs the QT interval.

The second overlap is the seizure threshold. A number of antiretrovirals, including darunavir, efavirenz, maraviroc, ritonavir, saquinavir and zidovudine, may increase the seizure risk associated with certain psychotropics, of which clozapine is the obvious example. Neither of these interactions is visible on a plasma-level check; both are predictable only from knowing what the two drugs do to the heart and to the cortex, which is exactly what makes the pharmacodynamic category easy to overlook and important to teach.

GOOD TO REMEMBER

Before adding citalopram or escitalopram to a QT-prolonging antiretroviral regimen, obtain a baseline ECG and monitor, because the two QT contributions add up even when neither drug alters the other's level. In the same spirit, when a seizure-threshold-lowering psychotropic such as clozapine is in the picture, remember that several antiretrovirals independently raise seizure risk, so the combined threshold is lower than either drug alone would suggest.

6.5 Choosing a mood stabiliser in HIV

Treatment of bipolar disorder in a patient with HIV is broadly similar to treatment in the general population, and the general pharmacotherapy belongs to the Mood disorders: pharmacotherapy notes. What changes is that the interaction burden reshapes the choice among the standard agents, so the selection becomes interaction-driven rather than efficacy-driven. Lithium is renally excreted, so cytochrome P450 interactions with antiretrovirals are unlikely; the catch is that its use is problematic in the renal impairment often seen in people living with HIV, where a narrow-index drug cleared by the kidney becomes difficult to hold in range. **Carbamazepine** is best avoided altogether, both because it carries significant drug interactions with antiretroviral therapy and because of its own risk of blood dyscrasias. Valproate is a known teratogen and is best avoided alongside other hepatotoxic drugs and in the presence of hepatic infection, a genuine consideration when viral hepatitis co-travels with HIV.

MOOD STABILISER	INTERACTION ISSUE IN HIV	PRACTICAL POSITION
Lithium	Renally excreted, so cytochrome P450 interactions with antiretrovirals are unlikely	Interaction-favourable, but problematic in the renal impairment often seen in people living with HIV
Carbamazepine	Significant interactions with antiretroviral therapy; risk of blood dyscrasias	Best avoided
Valproate	Known teratogen; hepatotoxicity concern	Best avoided with other hepatotoxic drugs and in hepatic infection

The logic reads cleanly off the table. The agent that interacts least through the liver, lithium, is limited by the organ that clears it rather than by the enzymes; the two anticonvulsant mood stabilisers are each constrained by a different hazard, carbamazepine by its interactions and marrow toxicity, valproate by teratogenicity and hepatic risk. Reasoning from those constraints, rather than from a memorised first-line list, is what the question is really testing.

6.6 The direct neuropsychiatric effects: efavirenz as the exemplar

The second half of the topic leaves interactions behind entirely: some antiretrovirals injure mood and cognition directly. **Efavirenz**, a non-nucleoside reverse transcriptase inhibitor, is the drug that defines the syndrome, and its recognised neuropsychiatric profile is broad: somnolence, insomnia, abnormal and vivid dreams, impaired concentration, depression, psychosis and suicidal ideation. The reassuring part is the time course, because these symptoms usually subside or diminish after **2-4 weeks** of treatment, so a patient warned in advance can often ride the early phase out. The caveats are two: subtler long-term neuropsychiatric effects may still occur once the acute phase settles, and the drug can exacerbate pre-existing psychiatric symptoms, which is why it is best avoided in a patient with a history of psychiatric illness.

The wider prescribing lesson generalises from that one agent. Psychiatric adverse events have been reported with most antiretroviral drugs, although for many the causal relationship remains uncertain and hard to separate from the psychological weight of the diagnosis itself. Efavirenz has been the most commonly implicated of them, and HIV guidelines accordingly suggest avoiding its use in patients with psychiatric comorbidity. The presence of suicidal ideation on the effect list is what lifts this above a tolerability footnote: it is a patient-safety issue, and it is the reason a psychiatric history is relevant to an antiretroviral choice at all.

> **PITFALL**
>
> Attributing new depression, insomnia or suicidal ideation in a patient recently started on efavirenz entirely to the psychological impact of learning an HIV diagnosis, and reaching for an antidepressant while leaving the antiretroviral untouched. The drug is the most commonly implicated agent for exactly these symptoms, and the early, largely self-limiting time course is a clue, not a coincidence. The regimen belongs in the differential for the psychiatric presentation before a primary disorder is assumed.

6.7 Efavirenz depression and psychosis, and the widening list of implicated agents

Looking more closely at efavirenz, of all the antiretroviral agents it is the one most often associated with depressive symptoms. Their course is variable and cannot be assumed: they remit spontaneously in some patients but persist in a few, so a wait-and-see stance does not serve everyone equally. Efavirenz has also been reported to cause frank psychosis in some patients, severe enough to necessitate withdrawal of the drug, and that is the crucial point about severity, because it means substitution rather than mere dose adjustment is sometimes the only safe course. More recently the list of implicated agents has widened beyond efavirenz alone: rilpivirine and efavirenz among the non-nucleoside reverse transcriptase inhibitors, and dolutegravir among the integrase inhibitors, have all been linked to depression.

That widening matters for how the topic is examined. The old teaching that efavirenz is the single neuropsychiatric offender is no longer the whole picture; the safer formulation is that efavirenz remains the principal and most consistently implicated agent, but that a newer non-nucleoside agent and a newer drug class have joined it. Keeping the three names together as a group is more useful than memorising efavirenz in isolation.

MNEMONIC

Efavirenz, Rilpivirine, Dolutegravir

The three antiretrovirals most consistently linked to depression, remembered as **E**favirenz and **R**ilpivirine among the non-nucleoside reverse transcriptase inhibitors, and **D**olutegravir among the integrase inhibitors. Efavirenz remains the principal offender; the other two mark where the concern has since spread.

6.8 Rilpivirine and the integrase inhibitors

Rilpivirine, the other neuropsychiatrically active non-nucleoside agent, produces depression, suicidality and sleep disturbances, a profile that resembles efavirenz's but carries a lower incidence of each event; it too is worth avoiding in a patient with a history of psychiatric illness. The **integrase strand transfer inhibitors** are the newer focus of concern. Dolutegravir, elvitegravir and raltegravir have each been associated with insomnia, depression and suicidal ideation, and the signal is strongest for dolutegravir, in which the symptoms characteristically appear within the first months of treatment.

AGENT OR CLASS	KEY NEUROPSYCHIATRIC EFFECTS	RELATIVE BURDEN AND TIMING	PRESCRIBING CAUTION
Efavirenz (NNRTI)	Somnolence, insomnia, vivid dreams, impaired concentration, depression, psychosis, suicidal ideation	Most commonly implicated; acute symptoms usually self-limit	Avoid in psychiatric comorbidity
Rilpivirine (NNRTI)	Depression, suicidality, sleep disturbance	Similar profile to efavirenz but lower incidence	Consider avoiding in a psychiatric history
Dolutegravir, elvitegravir, raltegravir (INSTI)	Insomnia, depression, suicidal ideation	Signal strongest for dolutegravir, within the first months	Monitor early, especially dolutegravir

The early timing of the integrase-inhibitor signal is a practical gift, because it makes attribution easier: neuropsychiatric symptoms that begin in the first months of a new integrase-based regimen point back to the drug more confidently than symptoms emerging years into stable therapy. Read together with the efavirenz story, the message across the whole class is consistent. Ask what the regimen is, ask when the symptoms started, and treat a close temporal link as a reason to reconsider the antiretroviral rather than simply to add a psychotropic on top of it.

- Before co-prescribing, check a maintained interactions resource, because the antiretroviral-psychotropic interactions are individually numerous and complex.

- Expect a boosted regimen to raise, and efavirenz to lower, the level of a shared substrate, and adjust the dose in the correct direction for the regimen in front of you.
- Take a full drug history: a potent inhibitor such as ritonavir can push a co-substrate, prescribed or recreational, into toxicity.
- Obtain a baseline ECG before adding citalopram or escitalopram to QT-prolonging antiretrovirals, and remember the shared effect on the seizure threshold with clozapine.
- Treat new depression, insomnia or suicidal ideation in a patient on efavirenz, rilpivirine or dolutegravir as possibly drug-related, not simply as a reaction to the illness.

In HIV, boosted regimens (ritonavir, cobicistat) raise psychotropic levels while efavirenz lowers them; efavirenz is also the antiretroviral most likely to cause depression, psychosis and suicidal ideation, so avoid it in anyone with psychiatric comorbidity.

Parkinson disease and its neuropsychiatry

7 · Parkinson disease psychosis: features and antipsychotic choice

A Parkinson patient who begins to see people who are not in the room has not acquired a second, independent psychiatric illness. Psychosis in Parkinson disease is a non-motor feature of the disease and of the drugs used to treat it, and it is common enough, and consequential enough, that recognising it and treating it safely is core neuropsychiatry rather than a footnote. The reason the marks concentrate here is a genuine therapeutic trap. The agents that reliably quieten psychosis in a younger patient are dopamine blockers, and dopamine blockade is precisely what the Parkinson brain cannot tolerate. Treat this psychosis the way you would treat schizophrenia and you will worsen the movement disorder you were asked to protect. This section defines **Parkinson disease psychosis**, works through its phenomenology and what drives it, sets out the anatomy that makes the wrong drug dangerous, and then follows the management logic: reduce the offending drug burden, treat any precipitant, and only then reach for a motor-sparing antipsychotic. The narrower and more dangerous problem of neuroleptic sensitivity in dementia with Lewy bodies is a different entity and belongs to the Dementias and neurocognitive disorders notes; it is named and pointed to here, not taught.

7.1 What Parkinson disease psychosis is, and how common it becomes

Parkinson disease psychosis is a non-motor symptom in which a patient with established Parkinson disease develops at least one of illusions, a **false sense of presence**, hallucinations or delusions, present for **at least 1 month**, once other causes have been excluded. That closing clause carries real clinical weight: a first psychotic episode in a Parkinson patient is a diagnosis of exclusion, and an intercurrent delirium, a metabolic disturbance or a newly added drug must be considered before the label is applied. The false sense of presence, the unshakeable feeling that someone is standing just beyond the edge of vision, together with fleeting illusions, is often the earliest and mildest expression and is easily dismissed by patient and family alike.

It is not a rare complication. Over the course of the illness a large proportion of patients are affected, and its importance is prognostic as much as symptomatic. Parkinson disease psychosis is associated with cognitive decline and, through it, with more frequent hospitalisation, heavier caregiver burden, earlier nursing-home placement and increased mortality. It behaves as a hinge point in the trajectory of the disease rather than as an incidental add-on, which is exactly why it deserves to be treated well and treated cautiously in equal measure.

up to 60% OF PD PATIENTS AFFECTED OVER THE DISEASE COURSE	**25-36 mg/day** MEAN CLOZAPINE DOSE FOR PSYCHOSIS IN PD	**40 mg/day** PIMAVANSERIN, PIVOTAL PDP TRIAL	**2016** PIMAVANSERIN FDA APPROVAL FOR PDP

7.2 The phenomenology: what the hallucinations look like

The characteristic disturbance is visual. Patients commonly experience **visual hallucinations** that take the form of complex, formed visions of people and animals, which fluctuate through the daytime and characteristically worsen at night when lighting is poor and cues are few. The content is frequently benign at first, familiar figures or small animals passing through the room, and the phenomenology tracks a recognisable arc as the disease advances.

Early in the course, and this is the discriminating point, the patient hallucinates while the sensorium is clear and while **retained insight** is preserved: the person knows the vision is not real even as they see it. That preserved insight, when the hallucinations are benign and non-distressing, is what can allow early psychosis to be managed conservatively, though only after an explicit assessment of safety and function, since insight alone does not establish that the patient is not at risk. As the disease progresses the picture darkens. The majority of patients go on to develop vivid hallucinations, disordered thinking and, often, paranoid ideation that tends to centre on the treating physicians and on family members. The loss of insight, rather than the appearance of hallucinations as such, marks the transition to a psychosis that changes what the clinician must do.

It helps to hold the spectrum in mind rather than a single symptom. The minor phenomena, illusions and the sense of a presence at the edge of vision, sit at one pole and are compatible with full insight and independent living; the formed, insight-poor, paranoid hallucinations sit at the other and drive the outcomes that make the syndrome matter. The nocturnal accentuation is itself a clinical clue, because it points to the contribution of impoverished sensory input and disturbed sleep and suggests correctable ground before any drug is changed. Grading the presentation along this axis, and re-grading it at review, is more useful than recording the mere presence or absence of hallucinations.

7.3 The anatomy that makes the treatment dangerous

The whole therapeutic dilemma is anatomical, so it is worth grounding. The motor syndrome of Parkinson disease, chiefly bradykinesia, rigidity and tremor, arises from the loss of dopaminergic neurons in the **substantia nigra pars compacta**. These nigrostriatal neurons degenerate slowly and lose tyrosine hydroxylase, and once approximately 60% of them are gone the pathway can no longer synthesise adequate dopamine and the clinical signs emerge. At post-mortem there is

depigmentation of the substantia nigra, the locus ceruleus and the vagal motor nuclei, and the surviving neurons contain Lewy bodies built around a core of alpha-synuclein, which is why the disease sits among the **synucleinopathy** group.

The pharmacological corollary follows directly. The projection from the nigra terminates in the **dorsal striatum**, and it is dopamine D2 receptor antagonism at that site that a dopamine-blocking antipsychotic produces. In a brain already dopamine-depleted, adding D2 blockade in the dorsal striatum removes what little motor drive remains. This is not an idiosyncratic adverse effect to be watched for; it is the predictable pharmacology of the drug meeting the pathology of the disease, and it is the single fact from which the entire antipsychotic-choice rule is derived.

7.4 What drives the psychosis and what precipitates it

Understanding the correlates changes management, because several of them are modifiable. Psychosis in Parkinson disease correlates most closely with dementia and with the duration and severity of the illness, and it also tracks with the number and dosage of dopaminergic medications and with the visible signs of excessive medication, particularly **dyskinesias**. A patient who is dyskinetic and hallucinating is, in effect, showing you two faces of the same dopaminergic excess. It is worth being precise about causation here: although Parkinson medications often precipitate the psychosis, they cause neither the comorbid depression nor the dementia that so often accompany the disease. The mood and anxiety complications of Parkinson disease, and the impulse control disorders and dopamine dysregulation syndrome that dopaminergic treatment can drive, are addressed separately and are only named here.

Beyond the disease itself, the additional risk factors are advanced age, sleep disturbance and visual impairment, each of which is worth screening for because each is at least partly correctable. The most actionable point of all is that intercurrent illness readily tips a stable patient into psychosis or frank delirium.

> **GOOD TO REMEMBER**
>
> Before reaching for any antipsychotic in a newly hallucinating Parkinson patient, actively look for a precipitant. An intercurrent pneumonia or urinary tract infection, dehydration, or a recently increased dopaminergic dose commonly triggers the psychosis, and correcting it may be all that is required. Uncorrected sensory impairment and disturbed sleep pull in the same direction.

7.5 First moves: rationalise the drug burden and treat the precipitant

Management is best thought of as a sequence, not a single prescribing decision, and most of it happens in the outpatient setting shared between neurology and psychiatry. Admission is reserved for the patient whose psychosis is driven by an intercurrent illness or delirium, or whose agitation and paranoia have become a safety risk, and the emergency presentation is the acutely agitated, frightened patient who cannot be contained at home. The acute target is not the hallucination itself but the driver behind it: treat the precipitating infection or metabolic upset, correct dehydration and sensory deprivation, and reduce the offending drugs before anything is added.

Drug rationalisation follows a deliberate logic rather than a rigid recipe. Because the antiparkinson agents differ in how much motor benefit they buy for the psychiatric cost they carry, the regimen is rationalised individually, withdrawing the least essential and most psychotogenic agents first, typically anticholinergics, then amantadine, MAO-B inhibitors and dopamine agonists, with COMT inhibitors and levodopa reduced last, before an antipsychotic is introduced at all. Levodopa comes down last precisely because it is the least dispensable for motor control. Only when the psychosis persists despite a rationalised regimen and a treated precipitant does the question become which antipsychotic, and by then a substantial fraction of patients no longer need one.

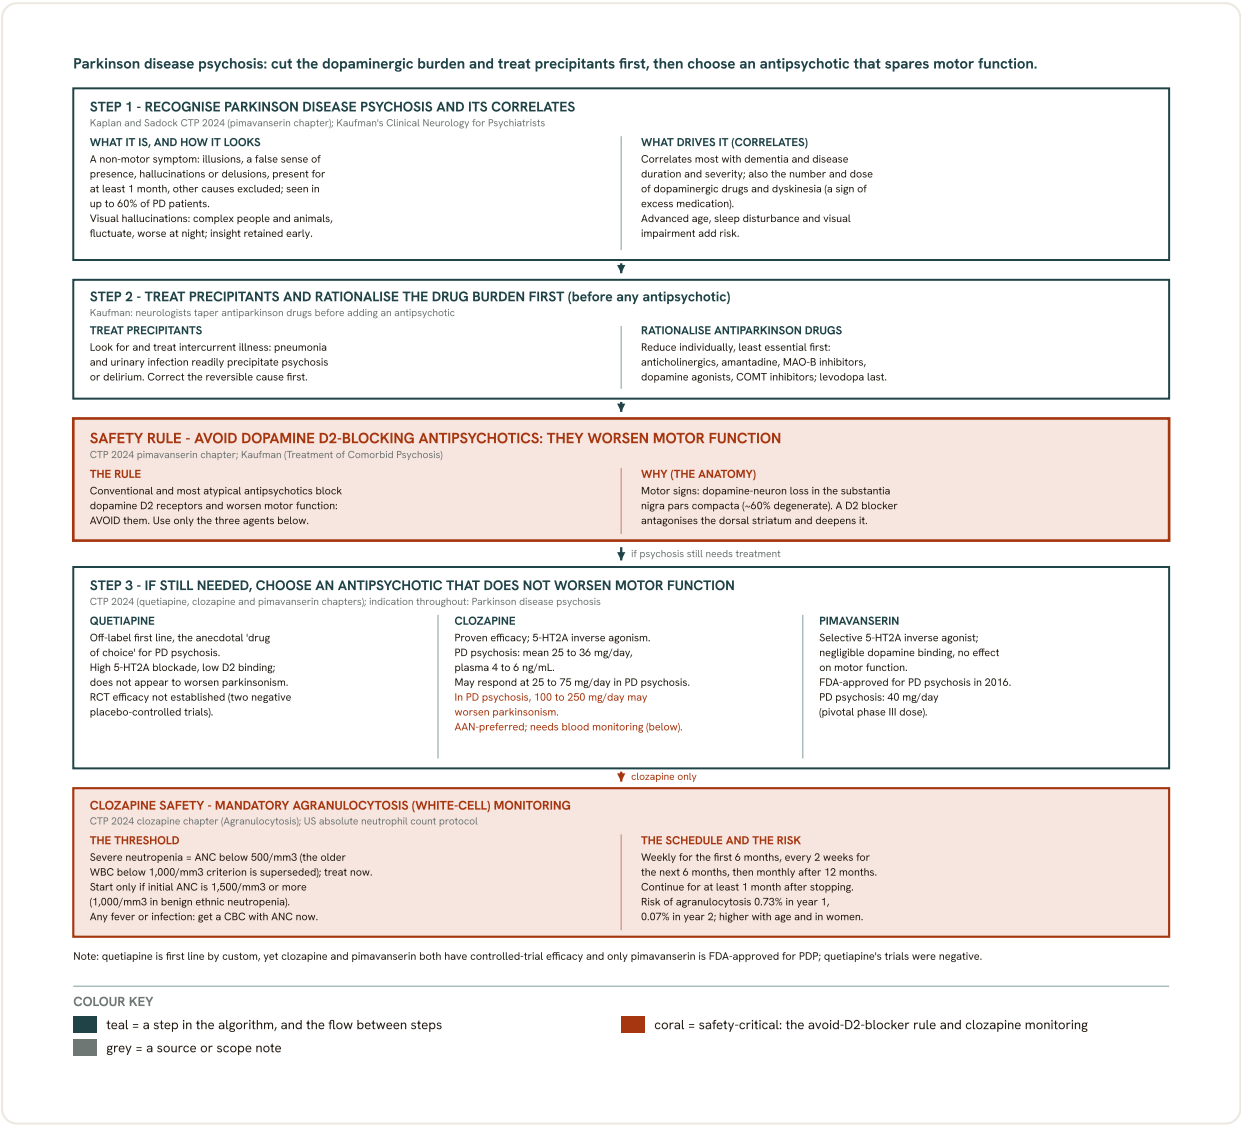


Fig 35.3 **Parkinson disease psychosis: the antipsychotic-choice algorithm.** Because dopaminergic drugs and disease severity drive the psychosis, the path first treats intercurrent precipitants and rationalises the antiparkinson burden, and only then reaches for a drug; the salient rule is that conventional and most atypical D2-blocking antipsychotics worsen motor function through antagonism in the dorsal striatum and are avoided, leaving quetiapine, clozapine with mandatory agranulocytosis monitoring, and the 5-HT2A inverse agonist pimavanserin as the motor-sparing choices. (Kaplan and Sadock 2024; Kaufman's Clinical Neurology for Psychiatrists.)

MNEMONIC

Anticholinergics, Dopamine agonists, Levodopa last

A simplified reminder of the core of an individualised taper before an antipsychotic is added: the least essential agents come down first (**A**nticholinergics, and also amantadine and MAO-B inhibitors), then the **D**opamine agonists, and **L**evodopa is spared till last as the agent the patient can least afford to lose for motor function.

7.6 The antipsychotic-choice rule and its safety boundary

This is the point on which the marks turn, and it is a safety point before it is an examination point. With the exception of low-dose clozapine and low-dose quetiapine, treating Parkinson disease psychosis with drugs that have antipsychotic actions frequently results in intolerable motor impairment, because those drugs act through dopamine D2 receptor antagonism in the dorsal striatum. The instruction that follows is unambiguous: avoid agents that block dopamine receptors, because they exacerbate the parkinsonism, and prefer agents whose antipsychotic action rests on serotonergic rather than dopaminergic blockade. The motor-sparing choices are quetiapine, clozapine and pimavanserin.

The logic of why the safe agents are safe is worth making explicit, because it is what the examiners are really testing. Ordinary antipsychotic effect is bought with dopamine D2 blockade, and in this disease that currency is exactly what the patient cannot spend. The motor-sparing agents instead lean on serotonergic action at the 5-HT2A receptor, which modulates the aberrant salience and perceptual disturbance of psychosis at a cortical level while leaving the already-fragile nigrostriatal dopamine transmission largely untouched. That dissociation, an antipsychotic action that does not depend on striatal dopamine antagonism, is the single pharmacological escape route the class exploits, and it is why a drug can quieten hallucinations without stiffening the patient. The corollary is that the more a candidate agent behaves like a classical D2 blocker, the less place it has here.

RULE **Do not use dopamine D2-blocking antipsychotics in Parkinson disease psychosis.** Conventional agents and most atypicals worsen the movement disorder through D2 antagonism in the dorsal striatum. The three motor-sparing options are quetiapine, clozapine and pimavanserin: pimavanserin is essentially non-dopaminergic, while low-dose clozapine and quetiapine retain some dopamine D2 antagonism but, at the low doses and strong serotonergic action used here, impose far less nigrostriatal motor burden.

AGENT	RECEPTOR BASIS	EFFECT ON PARKINSONISM	EVIDENCE AND REGULATORY STANDING
Quetiapine	High 5-HT2A blockade, relatively low D2 binding	Does not appear to worsen parkinsonism	Off-label first-line by convention; two placebo-controlled trials failed to establish efficacy
Clozapine	Serotonergic action; effective at low dose	Effective without worsening motor function at low dose	Established controlled-trial efficacy and tolerability; needs white-cell monitoring
Pimavanserin	Selective 5-HT2A inverse agonist, negligible dopamine binding	No effect on motor function	FDA-approved specifically for PD psychosis

The boundary of this rule must be drawn deliberately. What is described here is the treatment of Parkinson disease psychosis, which commonly coexists with Parkinson disease dementia; where it does, the antipsychotic risk-benefit balance shifts and the increased-mortality caution that applies to antipsychotics in dementia comes into play. The distinct and more hazardous phenomenon of neuroleptic sensitivity in dementia with Lewy bodies, and the antipsychotic-choice algorithm that follows from it, are owned by the Dementias and neurocognitive disorders notes and are not reproduced here. The two are easily confused and must not be run together.

PITFALL

Reaching for haloperidol or another high-potency typical antipsychotic when a Parkinson patient becomes acutely psychotic in a general ward. It is the intuitive move for agitation, and in this population it predictably precipitates severe motor deterioration through striatal D2 blockade. The safe reflex is the opposite of the general-psychiatry one.

In Parkinson disease psychosis, first rationalise the dopaminergic and anticholinergic burden and treat any precipitant; if an antipsychotic is still needed, choose a motor-sparing agent and never a dopamine D2 blocker.

7.7 Clozapine: low-dose efficacy at a haematological price

Clozapine is the reference agent, the one against which the others are measured, and its behaviour in Parkinson disease is instructive about the whole class. Used to treat the psychotic symptoms that are secondary to the dopaminergic agents themselves, patients with Parkinson disease psychosis may respond to relatively low dosages, in the region of 25 to 75 mg a day; higher dosages of **100 to 250 mg a day** may begin to exacerbate the parkinsonian symptoms. The efficacious range is narrow and low. Its ability to treat the psychosis at mean daily doses of 25 to 36 mg and at minute plasma levels of 4 to 6 ng/mL was historically important, because it pointed to 5-HT2A inverse agonism rather than dopamine blockade as the basis of the antipsychotic effect. A report from the American Academy of Neurology recommends clozapine as a preferred agent for psychosis in Parkinson disease.

The price is haematological and is the reason clozapine is not simply first-line for everyone. **Agranulocytosis**, defined by current monitoring as an absolute neutrophil count below 500 cells per mm3 (the older white-cell-count threshold of 1,000 cells per mm3 has been superseded by the neutrophil criterion), is a potentially fatal complication whose risk rises with age and is higher in women. The risk is front-loaded: it is **0.73%** during the first year of treatment and falls to 0.07% in the second year. Baseline and ongoing neutrophil-count monitoring is obtained according to current prescribing guidance, and although the mandatory United States clozapine registration requirement was withdrawn in 2025, any fever or sign of infection remains an immediate indication for a complete blood count with absolute neutrophil count.

- Initial absolute neutrophil count must be at least 1,500 per mm3, or at least 1,000 per mm3 in benign ethnic neutropenia, before the first dose.
- Count weekly for the first six months of treatment.
- Count every two weeks for the second six months.
- Count monthly thereafter, once the patient has passed twelve months.
- Continue monitoring for as long as the drug is taken and for at least one month after it is stopped.

7.8 Quetiapine, pimavanserin, and the trap of the single best drug

Quetiapine is, anecdotally, the antipsychotic drug of choice in practice for psychosis in Parkinson disease. It does not appear to worsen parkinsonism, and because of its high 5-HT2A blockade and relatively low D2 binding it has been used off-label as a first-line treatment. The awkward fact sitting underneath that popularity is that two placebo-controlled trials failed to show efficacy for quetiapine in treating the psychotic symptoms associated with the dopaminergic treatment of the disease, even though one study comparing it directly with clozapine found comparable levels of efficacy. It is convenient and well tolerated; its evidence base is weaker than its reputation.

Pimavanserin occupies a different position again. It is a selective **5-HT2A inverse agonist** and antagonist, acting at 5-HT2A and 5-HT2C receptors, and it is unique among drugs with antipsychotic actions in having negligible binding at dopamine receptors and no effect on motor function. That property is what took it to regulatory approval for Parkinson disease psychosis. The pivotal phase III trial, reported by Cummings and colleagues in 2014, used a uniform pimavanserin dose of 40 mg daily against placebo, begun after a 2-week lead-in, in 199 patients with Parkinson disease psychosis, and demonstrated significant improvement on the Scale for the Assessment of Positive Symptoms adapted for Parkinson disease compared with placebo. By the end of the sixth week, 37.2% of pimavanserin-treated and 27.8% of placebo-treated patients had responded, defined as a **50%** reduction on that scale.

TRAP **There is no single "best drug for Parkinson disease psychosis"; the right answer depends on which bar the question sets.** Quetiapine is the off-label drug of choice in practice, though its controlled trials were negative. Clozapine has the longest-established controlled-trial efficacy and tolerability, at the price of haematological monitoring. Pimavanserin also has positive controlled-trial efficacy and is the only agent approved by the FDA specifically for the indication. Preferred-in-practice, trial-proven and regulator-approved are different standards, and a stem that does not say which one it means is testing whether you conflate them.

8 · Depression and anxiety in Parkinson disease

Mood is not a bystander in Parkinson disease; depression is one of its defining non-motor features and, for the patient, often more disabling than the tremor. Parkinson disease carries a non-motor burden that is frequently heavier than its movement disorder, and depression sits at the centre of that burden. The marks cluster here for a reason that is clinical before it is academic: this depression is common, easy to overlook against a background of slowness and flat affect, and treated badly if the reflexes carried over from primary mood disorder are applied unchanged. Where the depression arises as a direct physiological consequence of the neurodegenerative process, the presentation is captured by the DSM-5 category of **Depressive Disorder Due to Another Medical Condition**; the broader nosology and staged pharmacotherapy of a primary depressive illness belong to the Mood disorders: pharmacotherapy notes and are not reproduced here. This section works through how common the depression is and why the figure is slippery, the recognition problem created by symptom overlap, the risk factors and the akinetic phenotype that carries them, the effect on cognition and quality of life, the companion problem of anxiety, the suicide risk that must be actively assessed, and finally a management logic that begins with the antiparkinson regimen rather than an antidepressant.

8.1 How common it is, and why the number refuses to settle

Reports of the prevalence of depression in Parkinson disease vary widely, because studies differ in how they define and ascertain a depressive disorder against the somatic background of the illness. That variability is not a footnote; it is the reason no single canonical figure can be quoted with confidence. By any measure, however, the prevalence is substantial. Clean prose sources commit only to a floor: studies typically report that **at least 30%** of patients with Parkinson disease manifest depression. Treat that as a lower bound rather than a headline. Part of the disagreement is methodological: instruments that count somatic items such as fatigue, slowing, and sleep disturbance tend to over-estimate depression in a disease that generates those signs for non-mood reasons, while instruments that discount the somatic items risk under-counting, and different studies sit at different points along that spectrum. The true proportion in any given clinic rises or falls with the diagnostic threshold applied and with whether minor and subsyndromal presentations are counted, which is exactly why apparently authoritative higher percentages should be held loosely unless the criteria behind them are stated. The clinically useful takeaway is not a precise percentage but the expectation itself: in any Parkinson clinic, a large minority of patients are depressed, and the clinician who is not actively looking will miss a treatable illness in a great many of them.

Akinesia	Assess risk	ECT
THE MOTOR SIGN LINKED TO DEPRESSION, NOT TREMOR	PD CARRIES INCREASED SUICIDAL IDEATION AND BEHAVIOUR	EFFECTIVE FOR BOTH MOOD AND MOTOR SYMPTOMS

8.2 The recognition problem: depression hiding inside the motor disorder

The single greatest obstacle to diagnosis is that the somatic vocabulary of depression and the phenotype of Parkinson disease speak over one another. Psychomotor slowing, reduced facial expression, a soft monotonous voice, disturbed sleep, fatigue, and reduced appetite are read by the somatic criteria of a depressive episode as evidence of depression, yet every one of them is also a face of the movement disorder itself. The masked, bradykinetic, hypophonic patient can look depressed while being euthymic, and the genuinely depressed patient can be dismissed as merely parkinsonian. Because the two conditions borrow each other's signs, the diagnosis has to lean on the cognitive and affective content that the motor disorder does not touch: pervasive low mood, anhedonia, hopelessness, excessive guilt, and a loss of interest that is out of proportion to the physical limitation.

A related discrimination is worth making explicit, because it changes treatment. **Apathy**, a flattening of motivation and goal-directed behaviour, occurs in Parkinson disease and is easily read as depression, but it can exist without dysphoria, guilt, or sadness, and it responds poorly to antidepressants that would help a true depressive episode. Separating a depressed patient from an apathetic one, and both from the involuntary flat affect of the disease, is a bedside skill that no rating scale substitutes for.

PITFALL

Attributing a Parkinson patient's slowness, poverty of expression, and disturbed sleep entirely to the movement disorder, and so never asking about mood at all. The overlap runs both ways: the same reflex also over-diagnoses depression in a patient whose flat, hypomimic presentation is purely motor. The corrective is to weigh the cognitive and affective content, low mood, anhedonia, guilt, hopelessness, rather than the somatic signs the two conditions share.

8.3 Risk factors and the akinetic phenotype

The risk factors are worth knowing precisely, because one of them is a classic examination discriminator. The most powerful predictors of depression in Parkinson disease cluster around illness severity and cognitive vulnerability rather than around tremor. Depression tracks the akinetic-rigid presentation, and not the tremor-dominant one, so that **akinesia** is a risk factor while tremor is not. That dissociation is the point examiners reach for, and it fits the wider pattern that the features signalling more diffuse and advanced disease are the ones that carry the mood risk.

- A previous history of depression, the strongest single predictor.
- Cognitive impairment.
- Akinesia, the akinetic-rigid phenotype, but not tremor.
- A young age at onset.

- A longer duration of illness.

Read together, these say that depression is more likely in the patient whose disease is cognitively involved, akinetic, early-onset, and long-standing. None of this is incidental to management: it tells the clinician which patients to screen hardest, and it warns that the appearance of depression may itself be a marker that the illness has moved into a more diffuse phase.

8.4 Why it matters: cognition, sleep, disability, and quality of life

Depression in Parkinson disease is not a parallel misfortune that leaves the neurological course untouched; it actively worsens it. When it occurs, the depression is associated with faster cognitive decline, disturbed sleep, and greater physical disability, so that a depressed Parkinson patient is functionally worse than the movement disorder alone would predict. The mechanism is partly a vicious loop: disrupted sleep worsens daytime motor performance and fatigue, disability erodes the very activity that ordinarily protects mood, and each increment of cognitive loss strips away some of the coping the patient had in reserve, so the depression and the disease amplify one another rather than merely coexisting. The relationship with cognition is bidirectional in practice: cognitive impairment predisposes to depression, and the depression, once present, predicts greater subsequent decline. The dementia of Parkinson disease itself, and its distinction from dementia with Lewy bodies and the one-year rule that separates them, are taught in this chapter's account of Parkinson disease dementia and in the Dementias and neurocognitive disorders notes; here it is enough to register that depression and cognitive deterioration pull each other downward.

The most striking single fact about its impact concerns quality of life. Of all the non-motor symptoms of Parkinson disease, depression is the one that correlates most closely with poor quality of life. That is a stronger claim than it first appears: it places mood, not the visible motor handicap and not the other non-motor complaints, at the head of what determines how a patient experiences the illness. It is also the strongest argument for treating the depression aggressively, because the lever it offers on overall wellbeing is larger than any other non-motor target.

8.5 Anxiety, the companion syndrome

Anxiety commonly travels with the depression, though it may also arise independently or fluctuate with the dopaminergic cycle. Where it accompanies the depression, the **anxiety** may present as another comorbidity in its own right, so that the two are often best thought of as a coupled disturbance rather than as entirely separate diagnoses, even though the anxiety can also stand on its own. Clinically the anxiety takes familiar forms, generalised apprehension, panic, and phobic avoidance, and in many patients it is at least as disabling as the low mood that drives it. It is worth eliciting specifically, because patients volunteer motor complaints far more readily than emotional ones, and an anxiety that is dismissed as understandable worry about a progressive illness is an anxiety that goes untreated.

Two practical points follow. First, because the anxiety is so often yoked to the depression, treating the depressive disorder frequently eases the anxiety alongside it, and anxiety that persists after the mood has lifted deserves its own attention. Second, the clinician should be alert

to anxiety that fluctuates with the patient's motor state through the day, since apprehension that rises and falls with the dopaminergic cycle points toward the timing of medication rather than toward a free-standing anxiety disorder. Where the anxiety is disabling in its own right, it is managed on the same principles as anxiety elsewhere, with the standing caveat that this elderly and frail population tolerates anticholinergic and sedative burden poorly, so agents and doses are chosen with that vulnerability in mind. The precise prevalence and subtype distribution of anxiety in Parkinson disease are not settled by the sources drawn on here, and the honest position is to name anxiety as a frequent, depression-linked comorbidity and to treat it on its clinical merits.

8.6 The suicide question: assess the risk, do not be reassured

Older teaching held that, although a chronic, progressive, disabling illness complicated by depression seems bound to carry a high suicide risk, earlier studies suggested Parkinson disease did not raise it, and that reassuring line was widely taught. Contemporary pooled evidence has overturned it: Parkinson disease is associated with increased suicidal ideation, suicidal behaviour and suicide mortality, and depression remains an important risk factor even though the absolute risk varies between studies. The practical consequence is the opposite of complacency: every patient, and especially every depressed patient, is assessed for suicidal ideation as a matter of course, with particular vigilance around a recent diagnosis, functional decline and changes in treatment.

One specific signal deserves to be held apart, because it is a different construct. Suicidality emerging in the context of deep brain stimulation is a recognised concern with its own literature, and it is addressed with that procedure in this chapter rather than folded into the natural history described here. Keeping the two apart, the raised baseline risk of the disease itself against the specific post-stimulation signal, is what allows the clinician to assess and counsel families accurately without either missing the day-to-day risk or overlooking the perioperative one.

TRAP **Relying on the older teaching that Parkinson disease does not raise the suicide rate.** That reassurance is outdated: contemporary evidence links Parkinson disease with increased suicidal ideation, behaviour and mortality, so a stem that offers "no increased suicide risk" as the expected answer is testing an obsolete fact. Assess suicide risk routinely, and remember that deep brain stimulation carries its own additional, separately recognised suicidality signal.

8.7 Treatment begins with the antiparkinson regimen, not the antidepressant

The first therapeutic move in Parkinson depression is counter-intuitive to a psychiatrist and it is the principle worth carrying away. Because the antiparkinson medicines can themselves improve mood, the antiparkinson regimen should be optimised before an antidepressant is added, so that some patients need no antidepressant once their dopaminergic therapy is adequate. This is not merely a sequencing nicety; it reflects that a component of the low mood is bound up with the state of dopaminergic transmission, and correcting the motor treatment can lift the mood as a by-product.

That dopaminergic logic extends to a specific and examinable agent. **Pramipexole**, a dopamine agonist used for both Parkinson disease and restless legs syndrome, has shown some promise in relieving the depression associated with Parkinson disease, and in a head-to-head comparison it proved superior to sertraline (Barone and colleagues, **2006**; the promise had been signalled earlier by Lemke and colleagues, 2006). The importance of this is conceptual as much as practical: a dopaminergic drug outperforming a standard serotonergic antidepressant reinforces the message that mood in this disease is partly a dopaminergic problem, and that optimising dopaminergic therapy is treatment for the depression and not only for the movement.

RULE **Optimise the antiparkinson regimen before prescribing an antidepressant.** The Parkinson medicines can improve mood in their own right, so an adequate dopaminergic regimen may lift the depression without a second drug, and a dopamine agonist has outperformed an SSRI for this indication. Only once the motor treatment is optimised does the question of a dedicated antidepressant properly arise.

8.8 Antidepressant choice, the serotonin-syndrome caveat, and ECT

When a dedicated antidepressant is needed, the historical and the modern evidence pull in different directions, and both matter. Earlier studies, corroborated in Tasman Psychiatry, found **tricyclic antidepressants** to be more effective than SSRIs for Parkinson-associated depression (Menza and colleagues, 2009). The difficulty is that these patients are generally elderly and therefore especially susceptible to the anticholinergic adverse effects of the tricyclics, and antidepressants as a class carry a greater risk of extrapyramidal side effects in this population. More recent research has demonstrated the efficacy of SSRIs here, which, set against the tolerability problem of the tricyclics in an elderly cohort, is why an SSRI is now a reasonable and often preferred first choice despite the older efficacy signal.

CLASS	EFFICACY SIGNAL	TOLERABILITY IN PARKINSON DISEASE
Tricyclic antidepressants	Older studies found them more effective than SSRIs, corroborated by Menza and colleagues	Elderly patients are susceptible to anticholinergic adverse effects; antidepressants also raise the risk of extrapyramidal side effects
SSRIs	More recent research has demonstrated efficacy in this population	Better tolerated in the elderly; the pragmatic first choice for many

A specific interaction generates more anxiety than it deserves. Prescribing an SSRI together with a **monoamine oxidase inhibitor** used in Parkinson disease, such as selegiline or rasagiline, can in theory precipitate **serotonin syndrome**. In practice the actual incidence is very low, because at the doses used to treat Parkinson disease selegiline selectively inhibits monoamine oxidase-B and largely spares the monoamine oxidase-A activity through which a dangerous serotonergic surge would occur. The clinical corollary is that the theoretical interaction should not by itself deprive a depressed Parkinson patient of an effective antidepressant, provided the usual vigilance for serotonergic toxicity is maintained.

GOOD TO REMEMBER

Do not reflexively withhold an SSRI from a depressed Parkinson patient merely because they take selegiline or rasagiline. At the low doses used in Parkinson disease the MAO-B selectivity keeps the risk of serotonin syndrome very low, so the combination is usually manageable with ordinary monitoring rather than an absolute contraindication.

The most robust option of all is one psychiatrists already know well. **Electroconvulsive therapy** is effective and safe for depression in Parkinson disease, and it has the notable bonus of temporarily improving rigidity and bradykinesia, an effect long recognised in the movement disorder alongside its antidepressant action. In an elderly, physically frail population in whom drug tolerability is the recurring constraint, a treatment that sidesteps the anticholinergic and extrapyramidal liabilities of the oral agents while acting on both mood and movement has an obvious and often under-used place. It is the treatment to remember for the severely or refractorily depressed Parkinson patient, precisely because it addresses mood and motor state together.

Treat the parkinsonism first, because dopaminergic therapy can lift mood on its own; if depression persists, favour an SSRI over the poorly tolerated tricyclic, do not fear the selegiline interaction at Parkinson doses, and reach for ECT when the depression is severe, remembering that the disease itself carries an increased suicide risk that must be actively assessed.

9 · Impulse control disorders and the dopamine dysregulation syndrome in Parkinson disease

The most damaging psychiatric complications of Parkinson disease are often the ones the treating team never thinks to ask about. A patient whose tremor and rigidity have been beautifully controlled may quietly gamble away the family savings, pursue sexual encounters wholly out of character, fill the house with things bought and never used, or spend hours a night dismantling and rebuilding the same small contraption. None of this is a coincidental second illness. Each is a behavioural complication of the dopaminergic drugs that restore movement, which is what makes the group at once dangerous, easy to miss and, encouragingly, often reversible. This section teaches three overlapping medication-driven syndromes: the **impulse control disorders**, the repetitive purposeless behaviour called punding, and the compulsive overuse of medication known as dopamine dysregulation syndrome. The mood and anxiety complications of Parkinson disease, and the psychiatric effects of deep brain stimulation in its own right, are close relatives that are addressed separately in these notes and are only named here.

9.1 The spectrum of impulse control disorders in Parkinson disease

An **impulse control disorder** in this setting is a failure to resist a drive towards a rewarding behaviour that has become excessive, repetitive and harmful. They are among the more dramatic

psychiatric complications of the disease, and they cluster into **four** cardinal appetitive behaviours that a candidate should be able to recite without hesitation:

- Compulsive or pathological gambling, from cards and casinos to online betting.
- Compulsive sexual behaviour, ranging from insistent demands on a partner to the pursuit of pornography or paid sex.
- Compulsive shopping, with impulsive and often ruinous buying.
- Compulsive or binge eating, frequently nocturnal and unrelated to hunger.

The behaviours are not rare curiosities. Taken together, one or more of them affects about **14%** of Parkinson disease patients in cross-sectional surveys, and a single patient may carry more than one at a time. What unites them is their appetitive, reward-seeking quality: the patient is driven towards a source of pleasure or relief and cannot moderate the pursuit, so the damage falls on finances, marriages and reputations long before it reaches the clinic. Recognising the pattern as a coherent syndrome, rather than as four unrelated misfortunes, is the first step towards treating it.

14%	four	3
OF PD PATIENTS DEVELOP ONE OR MORE IMPULSE CONTROL DISORDER	CORE BEHAVIOURS: GAMBLING, SEX, SHOPPING, EATING	FURTHER RISK FACTORS BEYOND THE DOPAMINE AGONIST

MNEMONIC

Bet, Bed, Buy, Binge

The four cardinal impulse control behaviours of Parkinson disease map cleanly onto four appetites: **Bet** for pathological gambling, **Bed** for compulsive sexual behaviour, **Buy** for compulsive shopping, and **Binge** for compulsive eating. All four are reward-driven behaviours, which is the feature that separates them from the aimless repetition of punding.

9.2 Why the dopaminergic drugs drive the behaviour

The mechanism is the reason these behaviours belong to neuropsychiatry rather than to primary addiction medicine. Parkinson disease depletes dopamine most severely in the motor circuits, while the reward-related dopaminergic projections are relatively spared until later, so a dose calibrated to rescue movement inevitably over-drives a reward system that was never as depleted to begin with. The replacement drugs flood the whole dopaminergic system, including the mesolimbic reward pathway that assigns value and salience to appetitive cues. Where the motor system is grateful for the extra transmitter, the reward system is over-stimulated, and behaviours that a healthy brain would weigh and abandon acquire an urgency the patient cannot override. The clinical fingerprint of this is that the offending drug and the disease are pulling in different directions in the same brain at the same moment, so the very act of dosing for good motor control tips a susceptible patient towards the behavioural complication.

The single most important fact for the examination follows directly from this. Treatment with a **dopamine agonist** is the most powerful risk factor for impulse control disorders in Parkinson disease, more so than the total dopaminergic load or levodopa alone. Agonists deliver a smoother,

more continuous stimulation of reward circuitry than the shorter-acting replacement of levodopa, and it is that sustained drive on the reward pathway that appears to matter. The corollary is clinically liberating: because the behaviour is pharmacological, it can usually be undone pharmacologically.

RULE **The impulse control disorders of Parkinson disease are iatrogenic and dopaminergic.** The dopamine agonist is the most powerful risk factor, and the same drug is therefore the treatment lever: reducing the agonist usually eliminates the behaviour, while deep brain stimulation, where appropriate, may help only indirectly by allowing the dopaminergic dose to be lowered. A behaviour you can switch off by adjusting a prescription is not a coincidental character flaw.

9.3 The risk profile: who develops an impulse control disorder

Not every patient exposed to a dopamine agonist develops an impulse control disorder, and the additional risk factors sketch a recognisable profile that is worth eliciting before and during agonist treatment. Beyond the agonist itself, the recognised associations are being unmarried, cigarette smoking, and a personal or family history of gambling. Each of these is easy to ask about and easy to record, and together they let the clinician flag the higher-risk patient before a drug is started rather than after the money has gone.

Two of these deserve a moment's thought because they point at the same underlying vulnerability. A personal or family history of gambling suggests a pre-existing sensitivity of the reward system, onto which the agonist then acts; cigarette smoking is itself a nicotine-reward behaviour and a marker of that same appetitive temperament. The practical message is not that these patients cannot receive agonists, but that the drug should be started with eyes open, with the risk named to the patient and the family, and with a plan to review the behaviour actively at each visit rather than waiting for a crisis to declare itself.

9.4 Recognition: the syndrome hides, so you must ask

The central clinical difficulty is that impulse control disorders are concealed. Gambling and sexual behaviour in particular carry shame, and the behaviour is frequently ego-syntonic, so the patient does not experience it as a symptom to report. The diagnosis is therefore made far more often on collateral history, when a spouse describes vanished savings, a secret credit card or a marriage under strain, than on spontaneous disclosure in the consulting room. A clinician who waits to be told will systematically miss the syndrome, and by the time the harm is loud enough to surface unprompted, the financial and relational damage is frequently irreversible. That asymmetry, a silent early phase followed by a catastrophic disclosure, is exactly why the syndrome rewards a screening habit rather than a reactive one.

The remedy is a low threshold and a direct question. Every patient on dopaminergic treatment, and especially every patient on a dopamine agonist, should be asked specifically and without euphemism about gambling, sexual behaviour, spending and eating, at the point of starting the drug and at every subsequent review, and the same questions should be put to the person who knows them best. A validated screening questionnaire for these behaviours in Parkinson disease

exists and can structure the enquiry, but the instrument is only as good as the willingness to raise the topic in the first place. The absence of a complaint is not the absence of the disorder.

GOOD TO REMEMBER

Ask about the four behaviours directly, and ask the family too. Because impulse control disorders are ego-syntonic and shameful, the single most useful move on a Monday morning is a routine, non-judgemental enquiry into gambling, sexual behaviour, spending and eating at every dopaminergic review, backed by collateral history. Screening tools help, but only after you have decided to ask.

9.5 Punding: repetition without reward

A distinct medication-induced complication is **punding**, defined as mindless, repetitive, purposeless behaviour. It is not appetitive and it is not pleasurable, which is precisely what separates it from the impulse control disorders. The classic vignettes are worth carrying: a patient who spends hours incessantly arranging the peas on a plate into small piles, who dismantles small constructions only to rebuild them, who opens and shuts a door over and over, or who folds and unfolds a newspaper without ever reading it. The acts are stereotyped, absorbing and utterly without product.

Three features define the syndrome and are examined together. Punding captures the patient's entire attention, displaces normal daily activities, and does not seem to yield any pleasure or excitement. That last point is the discriminator: the gambler is chasing a reward, but the punder is simply trapped in a loop. Related medication-driven behaviours sit on the same spectrum: **hobbyism**, which is repetitive activity pursued at a higher and more complex level than punding, and **hoarding**. The same aimless, stereotyped picture is seen outside Parkinson disease in autism and in amphetamine intoxication, a parallel that points squarely at dopaminergic and stimulant overactivity as the shared substrate.

COMPLICATION	CORE PHENOMENOLOGY	USUAL DRUG DRIVER	PLEASURE OR REWARD
Impulse control disorders	Gambling, sexual behaviour, shopping, eating	Dopamine agonist	Reward-seeking and appetitive
Punding	Aimless, stereotyped, purposeless acts	Dopaminergic medication	None; no pleasure or excitement
Hobbyism and hoarding	Higher-level repetition; accumulation	Dopaminergic medication	As for punding
Dopamine dysregulation syndrome	Compulsive overuse of medication	Levodopa	Craving, addiction-like

9.6 Dopamine dysregulation syndrome: addiction to the medication

The third syndrome shifts the compulsion from a behaviour to the drug itself. In **dopamine dysregulation syndrome**, once described as **hedonistic homeostatic dysregulation**, some Parkinson patients begin to overmedicate themselves. They express much greater concern for

their medication, usually levodopa, than their symptoms warrant, and describe their requirements in the language of obsession or addiction rather than of symptom relief. The medication has stopped being a means to move and has become an end in itself.

It is most useful to file this as a substance-addiction phenotype that happens to involve a prescribed antiparkinson drug. The patient escalates the dose beyond motor need, resists reduction, and hoards or hides supplies, and the behaviour tracks the short-acting, quickly rewarding forms of dopaminergic replacement rather than the smoother agonists. The three syndromes are not mutually exclusive, and a single patient can gamble, pund and overmedicate at once, because all three are expressions of the same dopaminergic overshoot read out through different circuits; recognising one should prompt an active search for the others. That distinction is the crux of the topic and the crux of the examination: the impulse control disorders are driven above all by the dopamine agonist, whereas dopamine dysregulation syndrome is a disorder of levodopa overuse. Two dopaminergic syndromes, two different culprits, two different management levers.

TRAP **Blaming the gambling on levodopa and the levodopa craving on the agonist.** Candidates lump all the dopaminergic behavioural complications together and attribute them to a single drug. The discriminator is precise: the impulse control disorders have the dopamine agonist as their most powerful driver, while dopamine dysregulation syndrome is the compulsive overuse of levodopa. Match the syndrome to the culprit and the management writes itself.

9.7 The withdrawal state and the danger of an abrupt taper

Because dopamine dysregulation syndrome is an addiction, it behaves like one when the supply is cut. Rapidly tapering or abruptly stopping dopaminergic medicines can throw the patient into a withdrawal state, with cravings, anxiety and drug-seeking behaviour, akin to an amphetamine user deprived of stimulants. This is not a psychological reluctance to lose a helpful drug; it is a recognisable pharmacological withdrawal, and it can be severe enough to defeat a well-intentioned attempt to bring the dose down.

The clinical consequence is that the drug cannot simply be withdrawn on the ward the way an offending agent might be stopped elsewhere. Reduction has to be gradual and negotiated, the withdrawal symptoms anticipated and managed, and the patient supported through a state that will otherwise drive them straight back to overuse. It is exactly the discipline of tapering any addictive agent, applied to a medication the patient also genuinely needs for movement, which is what makes dopamine dysregulation syndrome one of the harder problems in the whole field to manage well.

PITFALL

Yanking the dopaminergic dose down quickly once dopamine dysregulation syndrome is recognised, in the belief that stopping the misused drug is the obvious fix. An abrupt taper precipitates a withdrawal state of craving, anxiety and drug-seeking that both distresses the patient and undermines the reduction. The correct move is a slow, supported taper, not a sudden stop.

9.8 Management: working the dopaminergic trade-off

Management turns on a single tension that a figure can make vivid. Care is almost always delivered in the shared outpatient setting between neurology and psychiatry, with admission reserved for the rare patient whose behaviour has become an acute safety or financial catastrophe. The acute focus is to protect the patient and family from harm, chiefly by installing financial and relational safeguards and enlisting the family, while the longer-term focus is to bring the behaviour itself under control. In practice that acute containment can mean handing control of accounts and cards to a trusted relative, restricting access to online betting and shopping, and making the family a partner in monitoring, so that the behaviour cannot quietly resume between reviews while the slower work of adjusting the drug regimen proceeds. The definitive lever is pharmacological: reducing the dopamine agonist usually eliminates the impulse control behaviour. Deep brain stimulation, where it is appropriate, helps largely by allowing the dopaminergic drug burden to be lowered rather than by treating the behaviour directly, and it can occasionally leave impulsivity unchanged or worse; its own direct psychiatric effects are a separate matter taken up elsewhere in these notes.

The catch, and the reason this is a judgement rather than a formula, is that the drug you must reduce to stop the behaviour is the same drug that is controlling the movement. Every reduction that quietens the reward system risks re-emergent tremor, rigidity and slowness, so treatment is a deliberate weighing of motor benefit on one side against behavioural cost on the other, tuned to what each individual patient can tolerate.

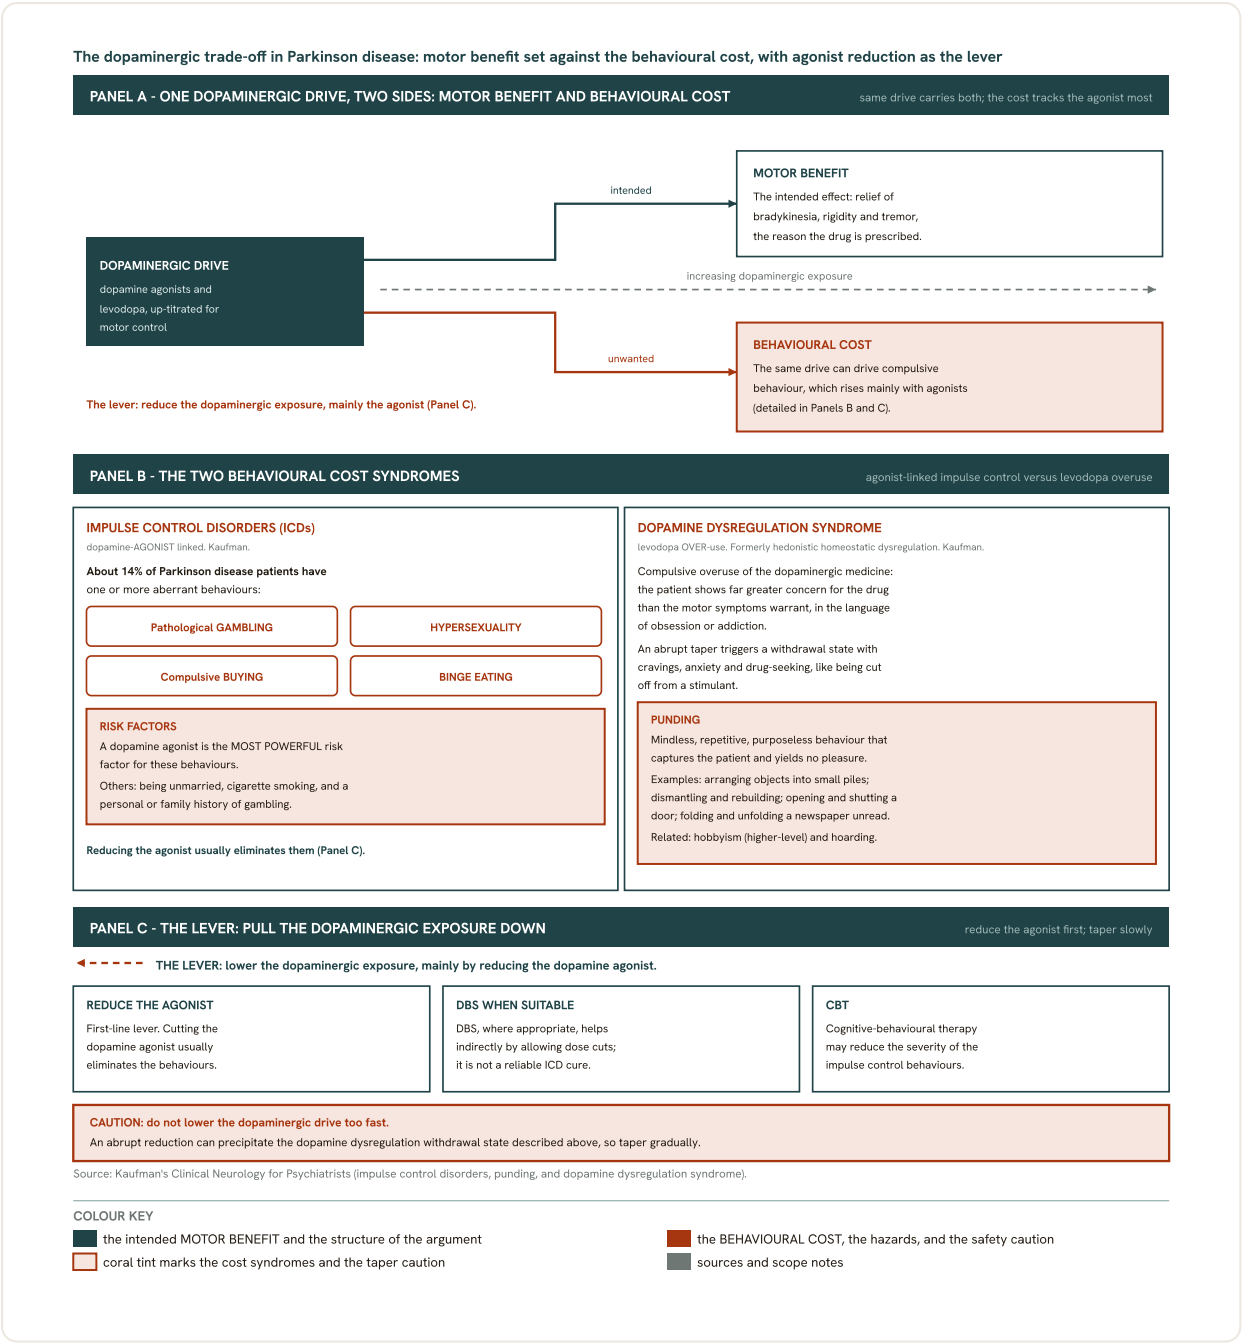


Fig 35.4 **The dopaminergic trade-off in Parkinson disease.** The same dopaminergic drive that relieves bradykinesia, rigidity and tremor also carries a behavioural cost: dopamine-agonist-linked impulse control disorders (pathological gambling, hypersexuality, compulsive buying and binge eating, affecting about 14% of patients) and the dopamine dysregulation syndrome of compulsive levodopa overuse, alongside the distinct repetitive behaviour of punding. The management lever is to lower the dopaminergic exposure, chiefly by reducing the agonist, supported by deep brain stimulation where appropriate and cognitive-behavioural therapy, while tapering slowly to avoid a withdrawal state. (Kaufman's Clinical Neurology for Psychiatrists.)

Alongside the drug adjustment, cognitive-behavioural therapy may reduce the severity of the impulse control behaviours and is a reasonable addition where it is available; the method itself belongs to the Psychotherapies, clinical assessment, MSE and nosology notes. Because the allowed evidence base here names no specific guideline, the safest counsel is to work from the reasoning above, screen actively, reduce the agonist first, and check current movement-disorder guidance for the detailed thresholds. The management of dopamine dysregulation syndrome differs in one crucial respect already made: its levodopa overuse must be tapered slowly to avoid

the withdrawal state, not reduced at the pace one might use for an agonist-driven impulse control disorder.

> The dopamine agonist is the most powerful driver of impulse control disorders in Parkinson disease, and reducing it or performing deep brain stimulation usually resolves them; dopamine dysregulation syndrome is instead the compulsive overuse of levodopa, and must be tapered slowly to avoid a withdrawal state.

10 · Cognitive impairment and Parkinson disease dementia

Cognitive decline is not an incidental extra in Parkinson disease; for a large share of patients it is where the illness eventually arrives. The motor picture of tremor, rigidity and bradykinesia is what brings the patient to the neurologist, but the cognitive syndrome that develops in established disease is what drives institutionalisation, caregiver burden and mortality, and it is squarely the psychiatrist's problem. The marks concentrate here for two reasons. The first is that the profile is distinctive: this is the textbook example of a subcortical dementia, and telling it apart from the cortical dementia of Alzheimer disease is a recurring examination task. The second is that its neurochemistry is counter-intuitive and clinically load-bearing. The dementia is not a dopamine problem, which is why decades of dopaminergic treatment do nothing for it, and recognising the prominent cholinergic deficit that underlies it is what points to the one class of drug that helps. This section works through how common the dementia is, what it looks like, how it is screened and tracked, how it is separated from dementia with Lewy bodies, and why its chemistry dictates its treatment.

10.1 How common the dementia is, and what predicts it

Dementia is a common endpoint of Parkinson disease rather than a rare complication, affecting a large share of patients over the course of the illness, and the risk is not evenly distributed. Its frequency rises in proportion to the patient's age, to the duration of the illness and to the burden of physical impairment, so the oldest patients with the longest-standing and most disabling disease are the ones in whom it should be actively anticipated. There is also a motor-subtype signal worth carrying: dementia is more frequent when the disease is dominated by akinesia and rigidity than when tremor predominates, so an akinetic-rigid presentation is a cognitive risk marker as well as a motor phenotype.

Timing matters for how the diagnosis is framed. Dementia typically develops later in the disease rather than in the early years, though mild cognitive impairment may already be present at or near diagnosis; a frank dementia that appears within months of the first motor signs should therefore prompt a rethink of the diagnosis rather than a label of Parkinson disease dementia. When the syndrome is established, the formal label under DSM-5 is **Major Neurocognitive Disorder Due to Parkinson's disease**, with the corresponding category sitting among the dementias in current international classifications. The stepwise predictors and the early-preservation rule together give a practical answer to whom to screen and when.

- Older age at the point of assessment.

- Longer duration of established movement disorder.
- Greater physical and motor impairment.
- An akinetic-rigid rather than tremor-dominant phenotype.

24-50%	2.3 points/yr	more than 1 year
OF PD PATIENTS AFFECTED; SHARE RISES WITH AGE AND DURATION	MEAN MMSE DECLINE IN PD (COHORT AVERAGE)	DEMENTIA AFTER PD DEFINES PDD, NOT DLB

10.2 The clinical profile: the archetypal subcortical dementia

What the dementia of Parkinson disease looks like matters as much as how common it is, because its shape is what separates it from the other dementias a psychiatrist meets. It is the classic example of a **subcortical dementia**, a useful shorthand more than a rigid category. The cognitive picture is dominated by inattention, poor memory and a characteristic difficulty in shifting mental sets, the executive inflexibility of a patient who cannot readily move from one rule or task to the next. Overlying all of it is **bradyphrenia**, a slowing of thought that is the cognitive counterpart of the bradykinesia seen in the limbs: the patient thinks slowly in the same way that they move slowly, and given time will often arrive at the right answer.

Two further features anchor the subcortical label. Gait impairment is almost invariable, so the dementia travels with a disorder of movement rather than appearing in an otherwise intact body, and language function is comparatively preserved, in contrast to the early linguistic breakdown of a cortical dementia. That combination, slow but not empty thinking, retrieval-type memory failure, executive rigidity, a disordered gait and spared language, is the signature to carry into the examination. It also has a practical corollary: because effortful retrieval and slowed processing dominate, a patient may perform far better when given time and cues than a brief bedside impression suggests, and cognition can be underestimated if it is assessed in a hurry.

MNEMONIC

Slow thinking, stuck sets, sound speech, shaky gait

The subcortical signature of Parkinson disease dementia: **slow thinking** is bradyphrenia, **stuck sets** is the difficulty shifting mental sets, **sound speech** is the preserved language function, and **shaky gait** is the almost invariable gait impairment, with inattention and poor memory completing the picture.

10.3 Subcortical against cortical: telling it from Alzheimer dementia

The single most examined discrimination is between the subcortical dementia of Parkinson disease and the cortical dementia of Alzheimer disease, because the two are predicted and managed differently and because the contrast crystallises what "subcortical" means. Parkinson disease dementia differs clinically from Alzheimer disease dementia along several axes at once, and it is more useful to hold the whole pattern than any single sign. The Alzheimer patient loses the memory trace itself and, before long, language; the Parkinson patient retrieves slowly from a memory that is comparatively better preserved and keeps their words. The table below sets the two profiles side by side.

DOMAIN	SUBCORTICAL PATTERN (PARKINSON DISEASE DEMENTIA)	CORTICAL PATTERN (ALZHEIMER TYPE)
Cognitive tempo	Bradyphrenia: markedly slowed thinking	Processing speed relatively preserved until late
Memory	Poor, inattentive, retrieval-type failure	Failure of encoding and storage, aided little by cueing
Set-shifting	Prominent difficulty shifting mental sets	Executive change tends to come later
Language	Comparatively preserved	Aphasic breakdown relatively early
Gait	Impairment almost invariable	Typically spared until late

The distinction is not merely academic. A retrieval failure that improves with cueing, cognition that is slowed rather than absent, and early gait disturbance should steer the clinician towards a subcortical process and away from an Alzheimer-type diagnosis, with consequences for prognosis, for what the family is told, and for which drugs are worth trying.

10.4 Screening and tracking the cognitive decline

Because the decline is gradual and the early years are usually spared, the practical question is which bedside instrument to use and how fast to expect change. Two brief cognitive tests are both valid for detecting cognitive impairment in Parkinson disease: the Mini-Mental State Examination and the **Montreal Cognitive Assessment (MoCA)**. They are not equivalent for this purpose. The MoCA is the more sensitive of the two, largely because it loads more heavily on the executive and attentional functions that fail first in a subcortical dementia, whereas the MMSE leans on the orientation and memory items that a Parkinson patient can pass until relatively late. For screening a patient at risk, the more sensitive instrument is the better first choice.

Rate of change is worth knowing for counselling and for spotting an accelerating course. On average, cohorts of patients with Parkinson disease lose about **2.3 points annually** on the MMSE, though individual rates vary widely with baseline cognition, education and phenotype, so the figure is a rough cohort benchmark rather than a precise individual expectation; a markedly steeper fall should still prompt a search for a superimposed process such as an intercurrent delirium or a second pathology. The detailed administration, scoring and cut-offs of these instruments are not reproduced here; they belong to the Psychotherapies, clinical assessment, MSE and nosology notes. What matters at the point of care is the choice of the more sensitive test and a realistic expectation of how quickly the numbers move.

GOOD TO REMEMBER

When you screen a Parkinson patient for cognitive impairment, reach for the more sensitive Montreal Cognitive Assessment rather than the MMSE, because the executive and attentional deficits of a subcortical dementia are exactly what the MMSE tends to miss. A normal MMSE does not reassure you that cognition is intact.

10.5 The one-year rule: separating it from dementia with Lewy bodies

Parkinson disease dementia and dementia with Lewy bodies share a pathology and, once both are advanced, share a clinical picture; what separates them at the bedside is nothing more principled than timing. The convention is an arbitrary **one-year rule**. Parkinson disease dementia is diagnosed when the dementia begins **more than 1 year** after a well-established parkinsonism, that is, when a patient who has had the movement disorder for years later becomes demented. Dementia with Lewy bodies is the mirror image: the dementia arrives before the parkinsonism, at the same time as it, or within a year of it, and some patients with the Lewy body dementia never develop parkinsonism at all.

The direction of that rule is the part candidates reverse under pressure, so it is worth fixing: an established movement disorder first with dementia later is Parkinson disease dementia; cognition first, or together with the movement disorder, is dementia with Lewy bodies. Once both dementia and parkinsonism are well established, the two seldom differ clinically, and it is reasonable to regard them as two ends of a single nosological continuum, sometimes grouped together as **Lewy body dementia**. The practical value of holding them as one continuum is that it stops the clinician forcing a false either-or at the bedside and instead brings the same vigilance to both: the same alertness to fluctuating cognition and to the perceptual disturbances that attend the Lewy body pathology, and the same caution before any drug that could destabilise an already fragile brain. The full diagnostic rule for dementia with Lewy bodies, its core clinical features and the neuroleptic sensitivity that governs its management belong to the Dementias and neurocognitive disorders notes and are not reproduced here; what a reader needs at this point is the timing discriminator and the recognition that the boundary is a convention rather than a biological threshold.

TRAP **Treating the one-year rule as a real biological boundary, and getting its direction backwards.** The split between Parkinson disease dementia and dementia with Lewy bodies is arbitrary, not a threshold with a mechanism behind it. Candidates lose the mark by reversing it: dementia appearing more than a year after an established movement disorder is Parkinson disease dementia, while dementia that precedes or coincides with the parkinsonism, or arrives within a year of it, points to dementia with Lewy bodies.

10.6 Why it is a cholinergic, not a dopaminergic, dementia

The most clinically load-bearing fact about this dementia is what does not cause it. Of the potential mechanisms, the evidence does not implicate dopamine deficiency: the simplest demonstration is that dopaminergic medicines, which transform the motor disorder, neither prevent nor relieve the dementia. The cognitive syndrome and the motor syndrome are driven by different lesions, and the drug that fixes one leaves the other untouched. This dissociation is the single idea from which the treatment of the dementia follows, and it explains why long-term levodopa does nothing to protect cognition.

What the evidence does point to is a cholinergic lesion. A likely contributor is an **acetylcholine deficiency**, and imaging supports it directly: positron emission tomography has shown a cortical cholinergic deficit in Parkinson patients with dementia that is more pronounced than that seen

in Alzheimer disease, itself the prototype of a cholinergic dementia. That the cholinergic loss should run deeper here than in the disease we most associate with cholinergic failure is a striking and examinable point. It reframes the dementia as a disorder of a different transmitter system from the one the motor disease has trained us to watch, and it makes a specific class of drug rational.

RULE **The dementia of Parkinson disease turns on a prominent cortical cholinergic deficit, not on dopamine deficiency.** Because dopaminergic medicines do not prevent or alleviate it, adding or increasing dopaminergic drugs is not a treatment for the cognitive decline; the cortical cholinergic deficit is the rational target, which is why a cholinesterase inhibitor, and not more dopamine, is the drug worth trying.

10.7 What follows for treatment, and what does not work

Management of the cognitive syndrome is mostly an outpatient, shared-care matter between neurology, psychiatry and the family, with admission reserved for the complications of advanced disease rather than for the dementia itself. The realistic targets are modest: slowing decline, protecting function and treating the neuropsychiatric features that ride alongside the cognitive loss, rather than reversing anything. It is worth being honest with families that the dementia is not reversible, and in particular that neither the dopaminergic medicines nor deep brain stimulation, whatever they do for the movement disorder, undo the cognitive decline; the psychiatric effects of the stimulation itself are a separate matter treated in their own right.

Where a drug can help, it follows from the chemistry. A **cholinesterase inhibitor** may provide modest symptomatic benefit for the cognitive impairment and, as a useful bonus, reduce the visual hallucinations that so often complicate the disease, so a single agent can address both the cognition and part of the psychosis that Parkinson disease otherwise obliges the clinician to treat with particular caution over antipsychotic choice. That dual action is worth pausing on, because it lets one rational choice serve two of the disease's hardest problems at once and spares the patient a second, riskier drug. Because no guideline is carried in the evidence here, the specific agent, dose and titration should be taken from current disease-specific guidance rather than from memory; the teaching point that survives is the class and its dual benefit.

It helps to frame the effort by phase. In the short term the aim is to establish the diagnosis honestly, set expectations, and remove the reversible aggravators that make any dementia look worse than it is. In the medium term it is to hold function with a cholinesterase inhibitor where it is tolerated and to keep the neuropsychiatric complications in check. In the long term the work shifts towards the family and towards planning for dependency, because this is a progressive disorder for which no available treatment alters the underlying trajectory. Across all of it, the non-drug measures carry as much weight as any prescription: caregiver education and support, treating a superimposed delirium promptly, and correcting the sensory and sleep problems that magnify confusion are the interventions that most often change how a patient actually does.

PITFALL

Reading the slowed, effortful cognition of a Parkinson patient as depression or as "just the bradykinesia", and either missing the dementia or escalating dopaminergic drugs in the hope of sharpening the mind. More dopamine does not treat this dementia and can precipitate hallucinations and confusion; the slowing to recognise is bradyphrenia, the cognitive counterpart of the motor slowing, not a low mood and not a dose that is merely too small.

10.8 Pulling the picture together

Set out in one line, Parkinson disease dementia is a late, subcortical, cholinergic dementia. It is late because cognition is usually intact for the first years of the movement disorder; it is subcortical because it slows thinking and disturbs gait while sparing language; and it is cholinergic because dopamine deficiency does not explain it and a cortical acetylcholine deficit does. Each of those three adjectives earns its place in an answer, and each carries a clinical consequence. The late onset is what the one-year rule uses to separate it from dementia with Lewy bodies, the subcortical profile is what distinguishes it from Alzheimer disease and tells you to screen with the more sensitive instrument, and the cholinergic basis is what makes a cholinesterase inhibitor, rather than more dopaminergic drug, the rational treatment.

Held together this way the topic is coherent rather than a list of facts: one pathology expressing itself late, in a recognisable cognitive shape, through a transmitter system that is not the one the motor disease taught us to watch. That is the frame to reproduce under examination pressure and to carry to the bedside.

Parkinson disease dementia is a late-onset, subcortical, cholinergic dementia: it appears more than a year after established parkinsonism, slows thinking while sparing language, and responds to a cholinesterase inhibitor and not to more dopamine.

11 · Deep brain stimulation and its psychiatric effects in Parkinson disease

Deep brain stimulation is where the psychiatrist meets the functional neurosurgeon. For the candidate this is a compact topic that carries marks out of proportion to its length, because it sits at the junction of movement-disorder neurology, functional neurosurgery and clinical psychiatry. The examiner is almost never testing operative technique. The questions turn on what the procedure can and cannot do for the parkinsonian patient, and on the mood and behavioural changes that follow it, some unwelcome and some beneficial. A psychiatrist is drawn in twice over: before surgery to help judge candidacy and to set expectations, and afterwards to recognise and manage the neuropsychiatric sequelae. The reasoning below is therefore squarely clinical rather than surgical, and it rewards a reader who can separate the motor promise of the procedure from its psychiatric fine print.

STN	GPi	Advanced PD
SUBTHALAMIC NUCLEUS, THE STANDARD MOTOR TARGET	GLOBUS PALLIDUS INTERNA, THE PALLIDAL ALTERNATIVE	THE ESTABLISHED INDICATION FOR STIMULATION

11.1 The procedure and where the electrodes sit

Deep brain stimulation is a functional neurosurgical procedure in which fine electrodes are implanted stereotactically into a deep grey-matter target and connected to an implanted pulse generator. In Parkinson disease the two established targets are the **subthalamic nucleus** and the **globus pallidus interna**, and the leads are tunnelled under the skin to a pacemaker-like device placed in the subcutaneous tissue of the chest, as shown in the accompanying figure. Once the system is switched on it delivers continuous high-frequency current, and because the stimulation is programmable and can be turned off, its effects are in principle adjustable and reversible in a way that a lesioning operation is not. The high frequency is the point: rather than destroying tissue, it modulates the pathological activity of the target and its connected motor network while the tissue itself is left intact, so a setting that later proves troublesome can simply be dialled back or off. This is the feature that makes the technique attractive in a progressive disease whose drug requirements keep changing.

The operation is not without hazard, and the candidate should be able to name its principal surgical risks:

- cerebral haemorrhage along the electrode track,
- infection of the implanted hardware, and
- fracture or failure of an electrode.

None of these is common, but each can be serious, and the risk of an intracranial bleed in particular is what keeps stimulation a treatment for carefully selected advanced disease rather than an early-stage convenience. Understanding the hardware also explains several of the psychiatric observations that follow, because mood and behaviour can shift both with the anatomical target chosen and with the way the current and the drugs are subsequently balanced against each other.

11.2 The motor payoff and the drug-sparing effect

The reason a patient accepts intracranial surgery is the motor benefit, and in advanced disease it can be striking. Stimulation of the subthalamic nucleus significantly relieves the motor symptoms of advanced disease, and the practical consequence for the patient is threefold. First, the involuntary writhing movements that chronic dopaminergic therapy eventually produces, the **dyskinesias**, are greatly reduced. Second, the abrupt swings between mobility and immobility, the **on-off fluctuations** that come to dominate late disease, become less severe. Third, because the stimulator now carries part of the therapeutic load, the patient can usually reduce the dopaminergic drug regimen while still maintaining mobility.

That drug-sparing effect is worth dwelling on, because it is at once a benefit and a hidden variable. A patient who was cycling through large, frequent doses of levodopa with distressing dyskinesias may, after stimulation, hold steadier on much less medication. The dyskinesias were largely a complication of the drugs themselves, so lowering the drugs while the stimulator maintains movement is precisely what abolishes them. But the same medication change alters the patient's dopaminergic tone, and dopaminergic tone is tightly bound up with mood, motivation

and reward. The motor win and the later mood story are therefore two faces of the same postoperative rebalancing, which is why a purely motor account of the procedure misses half of what the psychiatrist needs to know. Reading the operation as only a movement intervention is the commonest way to be caught out on the psychiatric half of the question.

11.3 The ceiling of benefit and the dopamine-resistant features

The single most examined idea here is that stimulation has a ceiling. It relieves the features of Parkinson disease that are dopaminergic and were levodopa-responsive, and it does little for the features that were never dopamine-driven: the **axial features** of gait and balance, and the non-motor burden of cognition and mood. One point cuts the other way and is worth banking: stimulation is not established to hasten cognitive deterioration, so the fear that an implant will accelerate a vulnerable patient into dementia is not borne out. That reassurance is narrow, though, because stimulation does not treat cognition and can still affect specific cognitive functions, so a cognitively vulnerable patient needs formal pre-operative cognitive assessment rather than blanket reassurance. The two-column summary opposite is worth memorising in exactly that shape, because the examiner's favourite error lives in the right-hand column.

PARKINSONIAN FEATURE	EFFECT OF SUBTHALAMIC OR PALLIDAL STIMULATION
Dyskinesias of chronic dopaminergic therapy	Greatly reduced
On-off motor fluctuations	Reduced
Overall dopaminergic drug burden	Usually lowered, with mobility maintained
Gait impairment	Not improved
Postural instability	Not improved
Cognitive impairment	Not improved, and not accelerated
Depression	Not a treatment target; mood may improve or worsen and needs monitoring

The reason for that ceiling is worth understanding rather than simply memorising. The motor features that respond are generated within the dopaminergic basal-ganglia loops that stimulation can reach and re-balance. Gait, balance, cognition and mood, by contrast, depend on circuits and transmitter systems that degenerate alongside the dopaminergic ones but lie outside the reach of a single deep electrode. A focal intervention on one node of one circuit cannot restore what is failing diffusely across several, which is why the ceiling is a property of the disease rather than a shortcoming of the technique, and why no refinement of the current will lift it.

RULE **Stimulation modulates the motor circuit; it does not cure the disease.** It re-balances an overactive output and lets the drugs come down, but the neurodegeneration continues and the dopamine-resistant features advance on their own timetable. Frame the procedure to patients as symptom control for a defined set of motor problems, never as a brake on the illness.

TRAP **Assuming stimulation improves every parkinsonian feature.** Candidates reason that if the technique helps the motor state it must also help the falls, the freezing of gait, the cognitive slowing and the low mood. It does not. Those are exactly the features that did not respond to levodopa, and stimulation obeys the same rule as the drug it substitutes for. A patient who is falling and cognitively declining will keep falling and declining after a technically perfect operation.

11.4 The neuropsychiatric footprint: mood, apathy and suicidality

Because the procedure rebalances dopaminergic circuits, it has a genuine psychiatric footprint, and this is the part of the topic the psychiatry examiner cares about most. In a study of parkinsonian patients undergoing stimulation, and depending on the target chosen, the postoperative dopaminergic drug regimen and other postoperative factors, some patients developed a new depression or had an existing one worsen, approximately **0.5% to 1.0%** had suicidal ideation or suicide attempts, and a larger proportion showed **apathy**. The reassuring counterpoint is that in many of these patients the disturbance was mild, transient, or amenable to treatment, so the finding argues for vigilance and structured follow-up rather than for withholding a procedure that can transform the motor state. Apathy deserves separate emphasis because it is quiet: it does not distress the patient in the way a low mood does, it is easily read by the family as laziness or depression, and it can persist even when the mood itself has recovered. At the bedside the discrimination that matters is between apathy and depression, because the two can look alike from across the room yet call for different responses: the apathetic patient has lost the drive to initiate but not necessarily the capacity for pleasure or the sense of hopelessness, and collapsing the two into one is how a treatable low mood gets missed, or an inert but not unhappy patient gets pushed with a drug that does little for drive.

MNEMONIC

Suicidality, Apathy, Depression

SAD is the postoperative neuropsychiatric watch-list to run after stimulation for Parkinson disease. **S**uicidal ideation or attempts in a small minority, **A**pathy in a larger share of patients, and **D**epression that is either new or a worsening of what was there before. Screen for the three, and remember that the presence of any of them is more often a signal to review the drugs and the stimulation than to abandon a working implant.

GOOD TO REMEMBER

A fall in mood, drive or motivation after stimulation is frequently mild, self-limiting or treatable. The instinct to switch off a stimulator that is controlling the motor state is usually wrong. The better first moves are to review the postoperative dopaminergic regimen, to treat the mood syndrome on its own merits, and to reassess before disturbing a system that is doing its motor job.

11.5 What shapes the psychiatric outcome

The direction of mood change after stimulation is not random. The same study frames the psychiatric outcome as contingent on three things: the anatomical target, the postoperative

dopaminergic medication regimen, and other postoperative factors. Each is a lever the clinician can think with. The target determines which neighbouring circuits the current spreads into. The medication regimen matters because stimulation typically allows the dopaminergic drugs to be cut, and a reduction in dopaminergic tone can of itself lower mood and drive; a patient whose motivation collapses after surgery may be tracking the drug change as much as the current. The third category, other postoperative factors, is a reminder that adjustment to a major operation, to a suddenly altered motor identity, and to shifting roles within the family all bear on how a patient feels once the novelty has passed.

This is also where the topic brushes against its neighbours without owning them. The impulse control disorders and the dopamine dysregulation syndrome that complicate dopaminergic therapy are a distinct behavioural problem taken up separately in these notes, and the low mood and anxiety of Parkinson disease as non-motor symptoms belong to that separate discussion too. What this topic owns is the narrower claim that the stimulation, the drug change it permits, and the perioperative course together set the psychiatric outcome, and that no single one of them explains it alone.

PITFALL

Attributing every postoperative mood change to the stimulator. Because the drug regimen is almost always reduced after implantation, a new depression or a new apathy may be driven by the falling dopaminergic dose rather than by the current. Reviewing, and where appropriate adjusting, the dopaminergic regimen is part of the workup of post-stimulation mood change, not an afterthought reached only once the stimulation settings have been blamed and found innocent.

11.6 From movement disorder to psychiatric tool

The psychiatric footprint of stimulation is not only a liability, and the other direction of the story is how the technique entered psychiatry at all. Deep brain stimulation is established in the treatment of Parkinson disease, where implantation of electrodes in the subthalamic nucleus significantly alleviates the motor symptoms of advanced disease. In the course of that work it was noticed that stimulation sometimes improved mood and **obsessive symptomatology** in parkinsonian patients. That incidental observation, set against a rapidly growing understanding of the neural circuitry that underlies psychiatric disorders, prompted formal studies of stimulation in patients with treatment-refractory depression and obsessive-compulsive disorder.

For the examination it is enough to hold the logic in sequence: a motor treatment produced an unexpected psychiatric dividend, circuit neuroscience supplied a rationale, and the technique was carried across into refractory psychiatric illness. The detailed management of treatment-refractory depression itself, and where a device intervention sits among its other options, belongs to the Mood disorders: pharmacotherapy notes, and this topic does not reproduce it. What Parkinson disease contributes to the wider story is the proof of principle that focal, adjustable stimulation of a deep target can move mood and obsessional symptoms, for better as much as for worse. For the parkinsonian patient in front of you, that same duality is the practical lesson: the current that steadies the motor state acts on circuits that also carry mood and drive, so the motor

and the psychiatric consequences of stimulation are never fully separable and are best reviewed together at every programming visit rather than handed off to different clinics.

11.7 Selecting the patient and the psychiatrist's role

Everything above converges on candidacy. The ideal candidate has advanced disease with disabling dyskinesias and on-off fluctuations that remain levodopa-responsive, because those are the features stimulation reliably helps. The features that predict a poor return are exactly the ones it does not touch: prominent gait and balance failure, and cognitive impairment. Pre-existing cognitive decline weighs twice, once because stimulation will not improve it and again because a cognitively frail patient tolerates the procedure and its later adjustments less well. Distinguishing Parkinson disease dementia, which is taken up separately in these notes, from dementia with Lewy bodies, whose diagnostic rule and neuroleptic sensitivity belong to the Dementias and neurocognitive disorders notes, is part of that cognitive assessment, and the standardised cognitive instruments used to make it are described in the Psychotherapies, clinical assessment, MSE and nosology notes.

The psychiatrist's contribution is concrete on both sides of the operation. Before surgery it is screening for a mood disorder and for suicide risk, and setting expectations the patient and family can actually hold. A documented pre-operative psychiatric baseline is quietly valuable here, because it is what later lets the team decide whether a postoperative change is genuinely new or simply the illness the patient already carried, and whether it is tracking the surgery, the drug reduction or the ordinary passage of time. Afterwards it is monitoring for the depression, apathy and suicidal ideation that can follow, and treating them without reflexively sacrificing the motor benefit that justified the surgery in the first place. The choice of antipsychotic in Parkinson disease psychosis is a separate problem again, addressed elsewhere in these notes and kept out of this one.

IN INDIA

At the movement-disorder clinics where stimulation is offered, the psychiatrist is often the clinician who frames the goal for a family weighing a major operation. The clinical decision to communicate is precise: stimulation targets the drug-refractory motor fluctuations and dyskinesias of advanced disease, and it does not treat the falls, the cognitive decline or the low mood that most erode independence. Screen mood and suicide risk before referral, and monitor them afterwards, as a built-in part of the pathway rather than something bolted on when a crisis has already appeared.

11.8 Pulling it together

Stimulation earns its place in the neuropsychiatry syllabus because it runs in two directions at once. It is a mechanical, adjustable intervention that reliably improves a defined slice of advanced Parkinson disease and reliably leaves the rest untouched, and it carries a psychiatric footprint that also runs both ways: capable of lowering mood, flattening motivation and, rarely, provoking suicidal thoughts, yet also capable of lifting mood and obsessional symptoms and, in that guise, of crossing into psychiatry as a treatment in its own right. Hold the motor ceiling and the

neuropsychiatric watch-list in mind together and the topic answers itself, in the examination and at the bedside alike.

Deep brain stimulation is symptomatic circuit modulation for advanced, drug-refractory Parkinson disease: it cuts dyskinesias, on-off fluctuations and drug burden, does little for the levodopa-resistant gait and postural instability, does not treat cognition or depression, and carries a small but real neuropsychiatric footprint of depression, apathy and rare suicidality.

Huntington disease and other movement disorders

12 · Huntington disease: genetics, CAG repeats, anticipation and trinucleotide pathology

Two facts decide almost everything in a Huntington pedigree: how long the CAG repeat is, and which parent the expanded allele was inherited through. Huntington disease is the model autosomal dominant neurodegeneration and the cleanest teaching example of a trinucleotide-repeat disorder, which is why examiners return to it so reliably. The genetics are deterministic in a way little else in psychiatry is, because a single expanded gene inherited from one affected parent will, given time, produce a defined disease with motor, psychiatric and cognitive components. This section builds the genetics from the molecular lesion upward, taking in the gene, the repeat, the penetrance bands, the instability that drives anticipation, and the relationship between repeat length and age at onset, and it marks the exact points where confident memory misleads. The clinical and cognitive syndrome, and the management of chorea, are taken up in the companion account of the psychiatric and cognitive features of Huntington disease and are only named here.

12.1 The disease and how it is inherited

Huntington disease, once called Huntington chorea, is an **autosomal dominant** neurodegenerative disorder characterised by midlife onset, a relentlessly progressive course, and a combination of motor, psychiatric and cognitive symptoms. Because the mutation is dominant, one copy of a fully expanded allele is sufficient to cause disease, and each child of an affected parent has an independent **50%** chance of inheriting the abnormal gene. There is no recessive carrier state, and a fully penetrant allele does not skip generations, although a reduced-penetrance allele can make a pedigree appear to skip one: a person who inherits a fully expanded allele and lives long enough will develop the illness.

The mutation lies in the **huntingtin (HTT) gene** on chromosome 4. Its cytogenetic address is 4p16.3, near the telomere of the short arm, and the lesion is an expansion of the triplet CAG within the gene rather than a point mutation or a deletion. This matters for how the disease behaves in a family, because a repeat can be counted, banded and tracked across generations in a way that a diffuse structural mutation cannot. The determinism of the inheritance is also what gives the disorder its particular weight in genetic counselling, a task handled separately in the account of predictive testing in Huntington disease.

~5/100,000	4.2 ± 0.8	0.57	± 1 triplet
HD PREVALENCE, UNITED STATES	MEAN CAG GAIN PER PATERNAL TRANSMISSION	R-SQUARED, REPEAT LENGTH VS AGE AT ONSET	LABORATORY ERROR AT THE BAND MARGINS

12.2 The molecular lesion: an expanded polyglutamine tract

The CAG repeat is not tucked away in an untranslated region; it sits in the coding sequence, in frame, so the expanded run of CAG triplets is transcribed and translated into an expanded run of glutamine residues in the huntingtin protein. Huntington disease is therefore one of a group of **13** neurodegenerative diseases caused by a **CAG trinucleotide repeat expansion**, where CAG denotes cytosine, adenine and guanine. Their shared final common pathway is a **polyglutamine** tract, and proteins carrying more than about **35 to 40** consecutive glutamine residues are thought to become neurotoxins, misfolding and aggregating in ways the neuron cannot clear.

This makes the mutation a toxic gain of function rather than a simple loss of the normal protein, and that distinction is worth holding onto because it explains the dominance. A recessive loss-of-function disorder needs both copies of a gene to fail before disease appears; here a single expanded allele produces a protein that actively harms the cell, so one copy is enough and the inheritance is dominant. The expanded huntingtin and the fragments cleaved from it accumulate and injure neurons, with the striatum among the most vulnerable regions, and the slow attrition of these neurons produces the chorea, the behavioural change and the cognitive decline that define the illness. The accompanying figure brings the gene, the repeat and the vulnerable striatum into a single view so that the molecular lesion and its target can be held together.

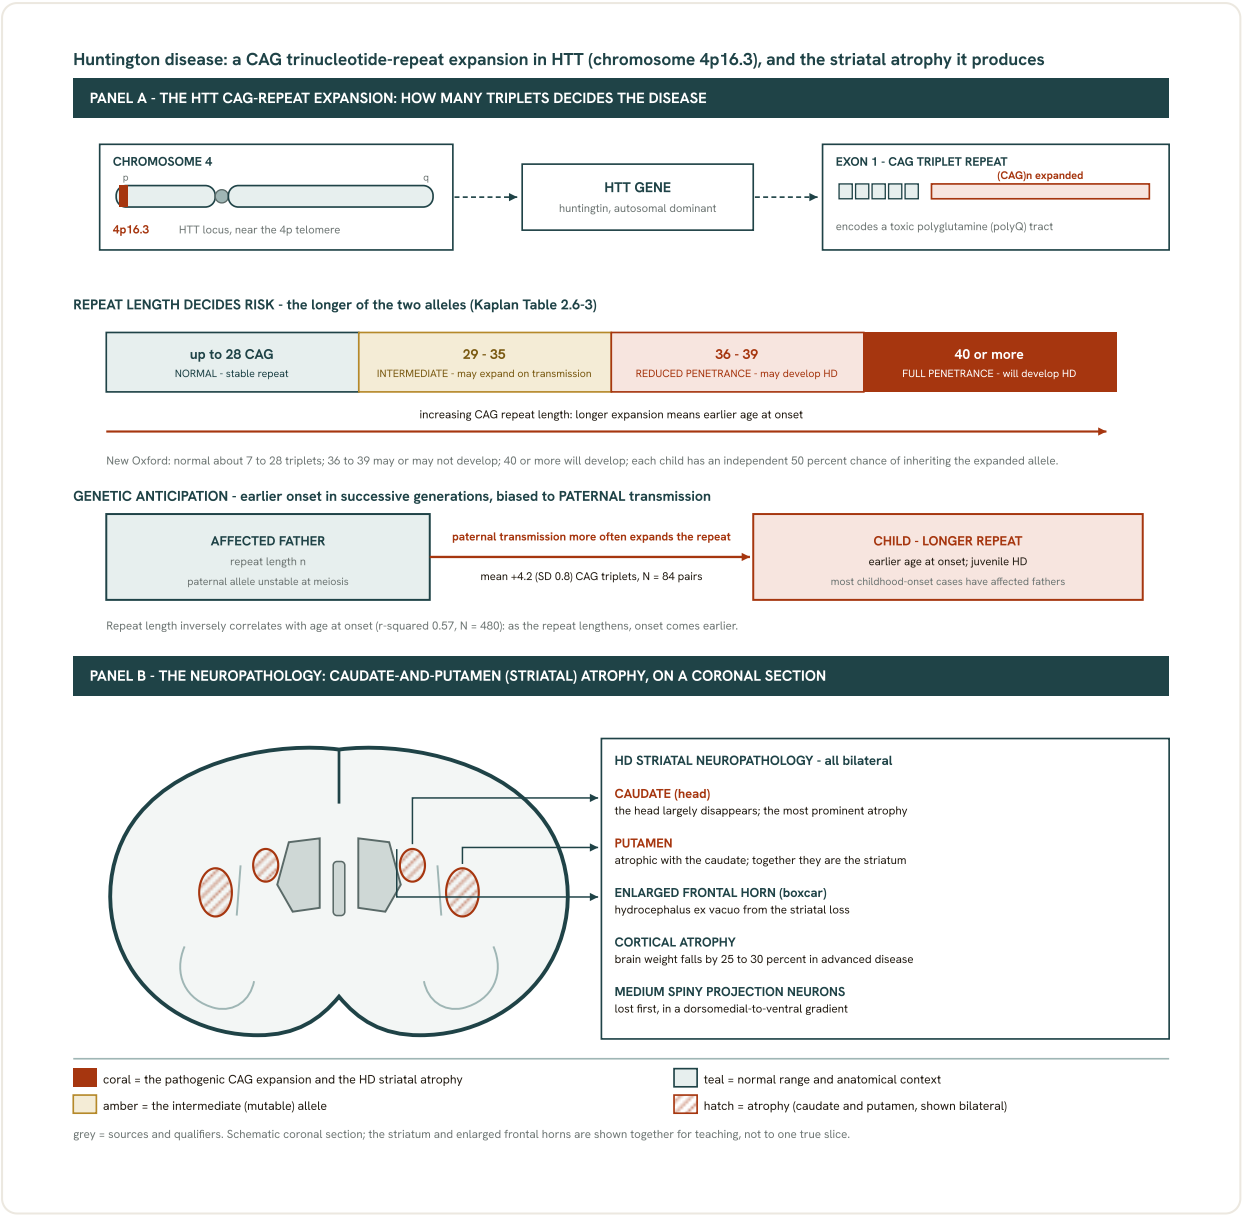


Fig 35.5 **Huntington disease genetics and neuropathology.** HD is an autosomal dominant disorder caused by a CAG trinucleotide-repeat expansion in the first exon of HTT at 4p16.3; the longer of the two alleles decides risk (up to 28 normal, 29 to 35 intermediate, 36 to 39 reduced penetrance, 40 or more full penetrance), and the repeat expands preferentially on paternal transmission, producing anticipation and juvenile-onset disease. The signature neuropathology is bilateral striatal (caudate and putamen) atrophy with enlarged frontal horns. (Kaplan and Sadock CTP 2024, and New Oxford Textbook of Psychiatry.)

The polyglutamine mechanism also explains two features that recur across this disease family and organise the rest of this section. The length of the repeat determines both whether disease appears at all and, if it does, roughly when; and the repeat is not faithfully copied from one generation to the next. Everything that follows, the penetrance bands, the phenomenon of anticipation and the inverse relationship with age at onset, is a downstream reading of those two properties of an unstable coding repeat.

RULE **An expanded CAG repeat becomes an expanded polyglutamine tract that poisons the neuron.** The repeat is in-frame and coding, so its length is read twice over: once as the number of glutamines that decide whether the protein is toxic, and again as the variable that sets penetrance and, inversely, the age at which the disease begins.

12.3 Reading the repeat length: the penetrance bands

Diagnostic genetic testing reports the number of CAG triplets in the longer of a person's two HTT alleles, and it is that count which is interpreted against a small number of bands. The bands are not arbitrary cut-points; they mark biologically meaningful transitions from an allele that is stable and harmless, through a mutable range that will not cause disease in the carrier but can expand in transmission, to alleles that carry **reduced penetrance** and finally to those with **full penetrance**, where the disease is certain in anyone who lives long enough. The interpretation table below follows the scheme set out in Kaplan and Sadock's Comprehensive Textbook of Psychiatry.

CAG REPEAT LENGTH	INTERPRETATION	PRACTICAL MEANING
Below 29	Not at risk; repeat length stable	Will not develop HD and does not transmit an expanding allele
29 to 35	Not at risk, though a few HD cases have been reported	Carrier is not expected to develop HD, though rare cases occur, and may pass an expanded allele to a child
36 to 39	May develop HD; risk rises with longer repeats	Reduced penetrance: disease may or may not appear in a lifetime
40 and greater	Will develop HD	Full penetrance; earlier onset with longer expansions

The single most examined transition is the step from the reduced-penetrance band to the fully penetrant one, because it is the point at which a result stops being probabilistic and becomes a certainty of eventual disease. It is worth being able to state the two break points without hesitation.

MNEMONIC

36 may, 40 will

The two break points that matter clinically: at a repeat length of 36 to 39 a person **may** develop Huntington disease, and at 40 and greater a person **will** develop it. Below that, the reassurance is qualified rather than absolute, which is exactly the ground the next section works over.

12.4 Where does normal end? The intermediate-allele band

The clean-looking table above hides a genuine disagreement between authoritative sources, and it is one examiners exploit. The trouble is not the disease-causing end of the scale, where the sources agree, but the lower boundary: where does the normal range stop, and what should the band just above it be called? The 29 to 35 range is the classic **intermediate allele**, a length that will not make its carrier ill but is unstable enough to lengthen into the pathological range when transmitted. Whether that band is filed under normal or set apart as a distinct category is where the sources part company.

TRAP **Assuming there is a single agreed number at which "normal" ends.** There is not. Kaplan and Sadock's Comprehensive Textbook of Psychiatry lists repeats below 29 and the 29 to 35 band as two separate rows, both classed as not at risk, whereas the New Oxford Textbook of Psychiatry puts the normal range at about 7 to 28 triplets and treats 29 to 35 as a distinct intermediate band. The sources agree on what decides who falls ill, reduced penetrance in the 36 to 39 band and full penetrance at 40 or more, but not on where normality stops or on whether 29 to 35 is named as healthy or as a mutable intermediate allele. Memorise a single boundary for "normal" and you will be caught precisely here.

The reason the disagreement is not merely academic is transmission. A person in the intermediate band is genuinely reassured about their own risk, yet the same instability that keeps them below the disease threshold is what can carry their child above it, so the counselling given to a carrier and to their offspring is not the same conversation. That asymmetry is the whole practical point of naming the band at all.

PITFALL

Reading a single borderline repeat count as an exact figure. Even the best laboratories carry an error of about plus or minus 1 triplet, so a result reported at the join between two bands cannot be treated as precise. A count of 35 against 36, or 39 against 40, sits inside the assay's own uncertainty, and such a result should be interpreted with that margin in mind rather than as a hard line cleanly crossed or not crossed.

12.5 Repeat instability and genetic anticipation

The repeat is **unstable at meiosis**: it is not copied faithfully from parent to child, and its length can change, usually by growing, when the gene is transmitted. That instability is markedly greater in the male germline, so the number of CAG triplets is more likely to increase when the gene passes through the father. The visible consequence in a family is **anticipation**, a fall in the age at onset across successive generations. Anticipation in Huntington disease is neither a statistical illusion nor an artefact of better detection; it follows mechanically from the combination of a strong correlation between repeat length and earlier onset and a tendency for the repeat to lengthen, most of all in paternal transmission.

The practical genetics of that instability reduce to a few points worth carrying:

- Repeat length can change between generations, so a child's count is not simply the parent's count read forward.
- The change is usually an increase, and increases are more likely in paternal transmission than in maternal transmission.
- Because expansion outweighs contraction, anticipation tends to run in a single direction down a pedigree, toward earlier onset.

A mother transmits the disorder with the same dominant probability, but the allele she passes is on average more stable, so maternally inherited disease tends to run truer to the parental age at

onset while paternally inherited disease is the one that can leap earlier and harder. That single asymmetry accounts for a large part of what makes Huntington pedigrees look the way they do.

GOOD TO REMEMBER

When you take a family history in a suspected Huntington case, ask which parent was affected and at what age they first presented. A pedigree in which each generation declares itself earlier than the last, particularly down a paternal line, is the clinical signature of an expanding repeat, and it should raise the possibility of a large expansion and an early, and in a child atypical, presentation before genetic results are back.

12.6 Repeat length and the age at onset

Repeat length predicts the age at onset, but only in part. Across large series the relationship is clear and monotonic: as the number of CAG repeats increases, the age at onset becomes earlier, so the longest expansions tend to declare themselves youngest and the shortest disease-causing alleles tend to present latest. This inverse relationship is one of the most reproducible quantitative findings in the disorder, and it is what underlies both anticipation and the juvenile presentations.

The crucial qualification is that the correlation, though strong, is far from complete. In the reference series relating repeat length to onset, repeat length accounted for rather more than half of the variance in age at onset and no more, which means a large share of the timing is set by other factors, among them modifying genes, environment and chance. The clinical translation is important and often examined. You can tell a person with a fully penetrant expansion that they will in time develop the disease, but you cannot read a precise age at onset off their repeat number, because two people carrying the same count can fall ill years apart. Repeat length bands the risk and shifts the expected decade; it is not a clock, and it should never be counselled as one. This is also why the shortest disease-causing alleles are the hardest to counsel: an allele just into the reduced-penetrance range may never declare itself in a normal lifespan, so the same length can mean lifelong health in one person and late disease in another, and honest counselling has to hold both possibilities open rather than collapse them into a single predicted age.

12.7 Juvenile onset and the transmission asymmetry

A minority of patients present in childhood or adolescence, and the genetics of this juvenile onset follow directly from everything above. It is the province of the largest expansions, and because those large expansions arise predominantly in the male germline, childhood onset is almost always inherited through an affected father. The chain from very long repeats, through paternal transmission, to the earliest onset is one of the most testable facts in the whole topic, and it ties the phenomenon of anticipation to its most extreme expression.

What the clean textbook sources withhold here is a number. They describe juvenile Huntington disease clinically and they tie it firmly to paternal transmission and to large expansions, but in these passages they do not state a specific CAG threshold that defines the juvenile or Westphal variant. It is tempting to supply the figure that circulates in teaching, and it is deliberately not asserted here, because the sources this account rests on do not carry it, and a remembered cut-off

is exactly the kind of confidently held number that turns out to be wrong. The distinct clinical picture of juvenile disease belongs to the companion account of the psychiatric and cognitive features of Huntington disease, and it is named here rather than taught.

12.8 How common it is, and the clinical bottom line

Huntington disease is uncommon but not rare, and its frequency varies markedly by population. In the United States the prevalence is about 5 cases per 100,000, with similar rates across most of Europe, while the New Oxford Textbook of Psychiatry gives a range of about 6 to 14 cases per 100,000 for North America, Western Europe and Australia. The disorder is much less common across Asia, roughly an order of magnitude lower in Japan and low in China, and it is also low in Finland and in Africa. Against this general pattern sit striking local concentrations, the best known being the communities around Lake Maracaibo in Venezuela, whose large affected pedigrees were central to mapping the gene.

For the practising clinician the genetics collapse into a compact frame. A patient with a suggestive movement disorder and a dominant family history in which onset marches earlier down the generations is very likely to be carrying an expanding CAG repeat, and the length of that repeat will place them in one of the penetrance bands and set the rough decade of their risk. Everything downstream, the counselling, the predictive-testing decision and the management of the psychiatric and cognitive syndrome, builds on that genetic core, and each is handled in its own companion account within these notes.

IN INDIA

Because Huntington disease is substantially less common across Asia than in Western populations, a young Indian patient presenting with chorea should not be labelled with an inherited polyglutamine disorder before acquired and treatable causes of chorea have been actively excluded. The lower pre-test probability of Huntington disease raises the relative weight of the reversible mimics, and when the diagnosis is genuinely suspected it is confirmatory genetic testing, not clinical impression, that should settle it.

Huntington disease is autosomal dominant: a CAG expansion of 40 or more triplets in HTT is fully penetrant, 36 to 39 carries reduced penetrance, the repeat lengthens and drives anticipation especially through the father, and longer repeats mean earlier onset.

13 · Predictive genetic testing and counselling in Huntington disease

Huntington disease is where predictive genetics first had to grow up. It is one of the few neuropsychiatric disorders with a **clearly defined genetic etiology** that can be identified in a person who is entirely well, years or decades before the first involuntary movement or the first change in mood. That single fact converts a family history into a testable question, and the question is not a laboratory formality: it asks a healthy adult whether they wish to know a diagnosis for which there is as yet no disease-modifying cure. A test run on a person who already has chorea and cognitive change merely confirms what the clinic already suspects; the whole

weight of this topic falls on the other use, testing a well relative who carries no signs at all, because there the result creates knowledge rather than explaining an illness that has already declared itself. The marks in this area sit almost entirely on the counselling around the test rather than on the assay itself. Examiners want to see that you understand who should be tested, how the result is interpreted, what it cannot tell you, and why the process is wrapped in mandatory counselling rather than handed over like a routine blood report.

1993	50%	40+	+/-1
DIRECT GENE TEST AVAILABLE SINCE	OFFSPRING RISK, EACH CHILD	CAG REPEATS: DISEASE CERTAIN	TRIPLET LABORATORY ERROR

13.1 A disorder built for the presymptomatic question

Ever since the causative gene was identified in **1993**, it has been possible to test directly for the Huntington repeat expansion using a simple PCR-based assay on DNA extracted from a blood sample, with testing offered through both commercial and academic centres. The molecular basis of the expansion, the phenomenon of anticipation and the instability of the repeat across generations belong to the account of the disease's genetics and are not the concern of the test itself; here the repeat length is simply the quantity the laboratory measures and the clinician interprets. What makes this disorder the teaching case for the whole of predictive medicine is the combination of a single, fully penetrant-at-threshold gene, an autosomal dominant pattern that puts identifiable relatives at high risk, and a long presymptomatic interval in which a well person can be given a near-certain answer about their future. That is a rare and uncomfortable position in clinical practice. Most of the conditions we screen for are probabilistic, modifiable, or both. Here the reader is dealing with a result that is close to deterministic and, for now, not modifiable, which is precisely why the ethical and counselling scaffolding around it is so heavy. A well person offered this test is being offered information rather than treatment, and information that cannot at present be acted on medically. That is why the right to decline, the right not to know, sits at the centre of the process and is protected as carefully as the right to be tested; a relative who decides they would rather live without the answer has made a legitimate choice, not a failure of nerve. The clinician neither steers the patient toward the certainty of a result nor toward the comfort of ignorance, and this deliberately even-handed stance is what marks predictive counselling out from most of medicine, where the clinician usually has a preferred direction. **Predictive testing** in this setting is therefore never framed as a laboratory decision. It is a considered choice made by an informed at-risk adult, and the clinician's task is to protect the quality of that choice rather than to encourage or discourage the test.

13.2 The fifty per cent rule and what a result means for a family

Huntington disease follows an **autosomal dominant** pattern, so a child of an affected parent carries a **50%** risk of developing the disorder. Each offspring, regardless of sex, has an independent one-in-two chance of inheriting the abnormal gene, and the word independent matters at the bedside: a family with three unaffected adult children has not "used up" its risk, and a sibling's negative result says nothing about the person in front of you. The genetic architecture also means the test is never truly about one individual. Because the abnormal allele

must have come from a parent, a positive predictive result in a young adult can retrospectively disclose the status of a parent who chose not to know, and a sibling's positive result, by confirming the gene is in the family, can raise the odds for relatives whose parental status was until then uncertain. This is the feature that raises the ethical issues the examiners expect you to name: a defined genetic disorder that can be diagnosed before symptoms creates real risks of discrimination in employment and insurance, and it forces questions about who else in the pedigree acquires knowledge they did not ask for. The discrimination risk is not hypothetical: a result sitting in a medical record can shape access to life cover or a post long before any symptom appears, which is one reason the decision to test is treated as private and considered rather than routine. The confidentiality problem is sharper still inside the family, because knowledge of one member's status leaks probabilistically onto the others. A young adult who tests positive has, in effect, said something about the parent from whom the gene descended, and a person who tests may reveal odds for siblings who wanted no part of the question. A counsellor cannot simply order the assay and report the number; the wider pedigree implications are worked through in advance, with the person helped to think about what they are willing to know and what they may inadvertently disclose. None of this is a reason to withhold testing from a competent adult who wants it. It is the reason the request is worked through slowly, with the wider family implications made explicit before any sample is drawn, so that the person understands they may be answering a question for relatives as well as for themselves.

13.3 Reading the repeat-length result

Interpretation rests on a small set of **repeat-length categories**, counted in CAG triplets, that translate a raw laboratory number into a counselling message. The categories used to interpret a predictive result are set out in the accompanying figure and in the table below.

CAG REPEAT LENGTH	INTERPRETATION FOR COUNSELLING
Under 29	Not at risk
29 to 35	Not at risk
36 to 39	Reduced penetrance: may or may not develop the disease
40 and greater	Will develop the disease

The clinically decisive band is the one in the middle. A repeat in the **reduced penetrance** range is not a normal result and not a diagnosis; it is a genuine uncertainty that has to be counselled as such. The lower bands are reported as not at risk of developing the disorder, while a result at the top of the scale carries the weight of a near-certain future diagnosis and must be delivered with the support that fact demands. The value of collapsing a continuous repeat count into these discrete messages is that patients do not act on triplet numbers; they act on whether they are, are not, or might one day become affected, and the categories are built to deliver exactly that judgement rather than a figure to be reinterpreted at home.

RULE **Under 36 repeats is not at risk, 36 to 39 is reduced penetrance where disease may or may not occur, and 40 or more means the disease will develop.** These three messages, not the raw triplet count, are what the patient leaves the room with. The counsellor's job is to convert a number on a report into the correct one of these statements and to make sure the intermediate band is never rounded up into a certainty or down into an all-clear.

13.4 The reduced-penetrance band is not a verdict

The single most common error in this topic, in the examination and in the clinic, is to treat a result in the reduced-penetrance range as though it were a diagnosis. The temptation is understandable: the patient came for a yes or a no, the number is above the not-at-risk threshold, and it feels honest to say the gene is present. But the categories are explicit that a repeat in this band means the disease may or may not develop, and a counsellor who collapses that into "you will get Huntington disease" has told the patient something the test does not support. The reverse error, reassuring the patient that a borderline expansion is nothing to worry about, is just as wrong. Both mistakes usually stem from a wish to give the family a clean answer where the biology does not provide one. The counsellor's task in this band is to hold the uncertainty open without inflating it into doom or dismissing it into nothing, to explain that a repeat in this range means the disease may occur and may not, and to help the patient plan around a probability rather than a promise. That is harder work than delivering an unambiguous positive or negative, and it is precisely where skilled counselling earns its keep; the honest answer here is a considered "we cannot be sure", which is a more difficult thing to say well than either a yes or a no.

TRAP **A result in the 36 to 39 band counselled as a certain future diagnosis.** Candidates who have memorised only "expansion equals disease" fall straight into it. The band carries reduced penetrance by definition, so the honest message is that the disorder may or may not occur, and the patient must be counselled to that uncertainty rather than to a verdict the laboratory has not delivered.

13.5 Laboratory limits and the borderline result

A predictive result is only as clean as the assay behind it, and the assay has a small but consequential margin. Even the best laboratories carry an error of +/-1 triplet, so a value sitting on the boundary of a category is inherently ambiguous: a report of 35 or 36, or of 39 or 40, may straddle the line between not at risk and reduced penetrance, or between reduced penetrance and certainty. A borderline number is therefore not something to interpret at face value; it is a signal to slow down, to discuss the uncertainty explicitly, and where necessary to confirm before any life-altering message is given. The safe habit is to treat any value within a triplet of a category boundary as provisional until it has been talked through and, where the stakes justify it, verified, rather than reading the printed number as though the assay were exact. A second, quieter trap lives in the raw report itself.

PITFALL

When the laboratory detects only a single allele of normal length, it is tempting to read this as a reassuring homozygous-normal result. In fact it usually means that both alleles happen to be the same length, and the finding must be interpreted with the laboratory rather than taken as excluding an expansion. Reading a single reported allele as an all-clear is a technical error that can hand a patient the wrong future.

13.6 Who asks to be tested, and who is counselled out

Uptake is far lower than the availability of the test would predict. Only a minority of at-risk individuals actually request predictive testing, and when they do the reasons are usually concrete rather than abstract curiosity. The commonest drivers cluster around a few life decisions:

- childbearing and reproductive planning
- long-term personal and career planning
- financial and insurance decisions

Just as important as who asks is who is counselled out. Those in an active major depressive episode are excluded from predictive testing until the depression is treated, because a result of either direction delivered into an active depressive episode carries an unacceptable psychological risk. This is not a bureaucratic hurdle; it is a clinical judgement that protects the patient from receiving a heavy answer at the point of least resilience. The pattern of who comes forward tells its own story. The people who request the test are usually facing a decision the answer would genuinely change, which is why reproductive and long-term planning dominate the stated reasons; abstract curiosity rarely survives the reality of an unmodifiable result, and many at-risk relatives who have thought hard about it choose to remain untested and live with the ambiguity. That self-selection is not an accident of the service but part of what keeps it safe, because the people who reach the sample are, on the whole, the people most prepared to use what it tells them.

GOOD TO REMEMBER

Screen for and treat major depression before offering premanifest testing, not after. Active major depression is an exclusion to the test, and stabilising mood first is what keeps the eventual result something the patient can absorb rather than something that lands on an unguarded moment.

13.7 The counselling protocol, before and after the test

Expert counselling is mandatory before **premanifest testing** is obtained, and the process is deliberately structured as **pre-test and post-test counselling** rather than a single consent conversation. The pre-test work ensures the person genuinely understands the potential psychological consequences of both a positive and a negative result and the real risks of discrimination in employment and insurance, and it checks that the decision is their own and not driven by a relative or an insurer. The post-test work carries the person through disclosure and the period that follows, whichever way the result falls, because the difficult phase is often not the

moment of the answer but the weeks after it. The reasoning behind the protocol is that a near-deterministic, presymptomatic genetic result is qualitatively different from an ordinary investigation and must be handled with matching care; the specific structure follows current predictive-testing guidance rather than any single fixed rule. The division into a before and an after is deliberate and worth understanding rather than memorising. The pre-test phase protects the quality of the decision, making sure the person has weighed a positive outcome against a negative one, understands the discrimination and family consequences, and is choosing freely rather than under pressure from a partner or an insurer. The post-test phase protects the person through the adjustment that follows the result, which for many is slower and harder than the moment of disclosure itself, and which is where an untested support gap does its damage. Consent for a predictive test is therefore not a signature captured on the day but a process that opens well before the sample is drawn and continues long after the answer is given. The pathway from referral through pre-test counselling to result disclosure and post-test support is summarised in the accompanying figure.

> No premanifest test in Huntington disease without structured pre-test and post-test counselling, and active major depression is a contraindication to testing.

13.8 Reproductive options and the relief paradox

For a person already known to carry the abnormal gene, the predictive question extends into the next generation, and two established options exist. **Prenatal diagnosis** early in pregnancy allows the fetus to be tested, and **preimplantation genetic testing** of embryos generated by IVF allows unaffected embryos to be selected before a pregnancy is established. Both keep the choice about transmission in the hands of the carrier rather than leaving it to chance, and both belong in the counselling conversation for anyone whose result is positive and who is thinking about children. The chapter closes on a finding that surprises trainees. Given how heavy a positive answer is, one might expect testing to leave most people worse off, yet the opposite is generally observed: most patients feel relieved to learn their status, whether that status is positive or negative. The likely explanation is not that the news is easy but that the cohort is selected. The people who reach the test are those who have chosen it and worked through what it means, and those in an active depressive episode have been counselled out, so the group that actually receives results is unusually prepared for them. Put another way, the reassurance seen at follow-up is bought by the front-end work: the counselling that prepares people for either answer and the exclusion that keeps the most vulnerable out of the queue until they are well. Strip either safeguard away and the same result would land very differently, which is why the finding is so often misread. It does not show that the news is easy to bear; it shows that a well-run process can make even hard news bearable. That relief is a product of the protocol, not a reason to relax it, and it is the strongest practical argument for keeping the counselling around the test as rigorous as the test is simple.

14 · The psychiatric and cognitive features of Huntington disease and their management

Huntington disease is where the psychiatrist meets a movement disorder that is also a dementia and a mood disorder, and often meets it before the neurologist has made the diagnosis. The chorea is what names the illness, but for the patient and the family the psychiatric and cognitive changes are frequently the earlier and the more disabling burden, and they are where a large share of the clinical management sits. Marks concentrate on this material for three reasons that are worth holding in mind as the topic unfolds. The suicide risk is genuinely high and it is high early, before motor disease and sometimes before the person even knows their risk. The cognitive decline is the textbook example of a subcortical dementia and is regularly contrasted with the cortical picture of Alzheimer disease. And the drugs that most reliably control the chorea are the same drugs that can deepen depression and precipitate suicidality, which turns the management into a safety problem before it is a symptom-control problem. The genetics of the trinucleotide repeat expansion, and the predictive testing and counselling that flow from it, are taken up separately in this chapter and are named rather than re-taught here; this section works from the neuropathology through the psychiatric spectrum, the suicide risk, the cognitive profile and, finally, the management of the chorea and its psychiatric hazards.

14.1 The neuropathology that shapes the whole syndrome

The clinical picture reads directly off the pathology, so it is worth grounding first. The most prominent change is bilateral striatal atrophy: the caudate and the putamen waste, the head of the caudate largely disappears, and the frontal horn of the lateral ventricle enlarges to fill the space it leaves, producing the **hydrocephalus ex vacuo** that is visible on MRI and on gross specimen alike. The picture is one of an atrophic striatum with a dilated ventricle and a thinned cortex. The cerebral cortex itself becomes markedly atrophied, while other basal ganglia structures, particularly the globus pallidus and the subthalamic nucleus, are spared relatively more, and in advanced disease the brain may lose **25% to 30%** of its weight.

Within the striatum the cell loss is selective and ordered. It falls chiefly on the **medium spiny projection neurons**, in a gradient with the earliest and most severe loss in the dorsomedial striatum, while in the cortex it is the large neurons of the deeper laminae, layers III, IV and V, that are depleted. This anatomy explains the triad. Degeneration of the striatal output disinhibits movement and generates the chorea; disruption of the parallel frontostriatal circuits that loop through the same territory produces the executive slowing, the apathy and the personality change; and the diffuse cortical involvement underlies the later global cognitive decline. The psychiatric and cognitive features are therefore not incidental comorbidity but expressions of the same circuit pathology that drives the movement disorder.

14.2 The psychiatric spectrum and how often it appears

Psychiatric disturbance in Huntington disease is common, disabling and, importantly, not merely a psychological reaction to a fatal diagnosis. Over the course of the illness the great majority of patients develop some form of noncognitive psychiatric disorder, typically within 10 to 15 years of motor onset, and a substantial minority, up to 20%, present to services with psychiatric

symptoms in the first instance. The syndromes span the whole of psychopathology. Depression dominates; beyond those meeting criteria for a major depressive episode, a further 10% to 20% carry a different, non-major depressive syndrome. Psychotic phenomena occur across the illness, and a minority develop mania, which is easily mistaken for the disinhibition of the frontostriatal syndrome rather than recognised as an affective switch. The distinction is not academic: a sustained, mood-congruent elevation with reduced sleep and grandiosity calls for mood stabilisation, whereas the flat, undirected disinhibition of circuit damage does not, and treating one as the other wastes time and risks harm.

up to 80%	11% to 13%	30% to 40%	up to 10%
DEVELOP A NONCOGNITIVE PSYCHIATRIC DISORDER DURING THE ILLNESS	LIFETIME PREVALENCE OF PSYCHOTIC PHENOMENA	DEVELOP MAJOR DEPRESSION	DEVELOP MANIA

The practical consequence is that a patient with Huntington disease may need treatment for a mood disorder, a psychosis or both, and the primary nosology and pharmacotherapy of those syndromes are not rewritten by the neurological diagnosis. The general management of depression belongs to the Mood disorders: pharmacotherapy notes, and the general management of psychosis to the Schizophrenia: management, antipsychotics and adverse effects notes; what changes here is the safety framing, which the sections that follow develop.

14.3 Apathy, irritability and the personality change that comes first

If depression is the most frequent syndrome, personality change is the most characteristic, occurring in up to 50% of patients by the middle of the disease course. The two features that dominate this change, and that families find most troublesome, are **apathy** and irritability, and their importance to the examiner is that they may appear before there is any overt movement disorder at all. A patient who is gene-positive but pre-manifest, or in whom the chorea is still subtle, may present to psychiatry with a slow erosion of drive, initiative and emotional engagement long before a neurologist would call the disease clinically manifest.

Apathy here should not be collapsed into depression. It is a loss of motivation and goal-directed behaviour that can occur with a relatively preserved, non-dysphoric mood, and it tracks the frontostriatal circuit disruption described earlier rather than a primary affective disturbance. Distinguishing the two matters because the apathetic patient will not respond to, and may be worsened by, the sedating strategies that suit an agitated depression, whereas the irritability, which can shade into aggression and impulsivity, may itself demand containment. Because these changes can predate the movement disorder, they are frequently the reason such a patient is first brought for neurological and diagnostic assessment; the separate question of predictive testing in still-asymptomatic relatives is taken up separately in this chapter.

14.4 Suicide: the safety-critical feature

No fact in this topic is more important to carry into a clinic than the suicide risk. It is much higher than in the general population, it is present across the illness, and it is present in the prodrome, sometimes before the patient is even aware of being at risk. Lifetime suicidal ideation

has been estimated at 20%, and the lifetime rate of completed suicide at 6.6%. Prospective data from the ENROLL-HD cohort put the incident rate of suicide among gene carriers at 72 per 100,000 person-years and of suicide attempts at 306 to 375 per 100,000 person-years, figures that run at approximately **10 times** the rate in non-carriers.

Two clinical implications follow. First, suicide screening cannot be confined to the depressed or to the advanced patient; it belongs at every contact, including the pre-manifest phase and the period around genetic testing, when the disclosure of risk is itself a hazard. Second, the choice of antidepressant is shaped by the risk. Because overdose is a real danger in this population, antidepressants that are less toxic in overdose are preferred.

RULE **Screen for suicide at every stage of Huntington disease, including the prodrome.** The elevated risk begins before motor onset and before diagnosis, and it can be driven by the disclosure of genetic status rather than by symptom burden, so a screen that waits for depression or for clinical manifestation is a screen that comes too late.

GOOD TO REMEMBER

When treating depression in Huntington disease, the raised suicide risk pushes the choice toward selective serotonin reuptake inhibitors, which are less toxic in overdose than the tricyclic antidepressants or the monoamine oxidase inhibitors. The reasoning is a safety one and applies before any question of comparative efficacy.

14.5 The cognitive profile: a subcortical dementia

The dementia of Huntington disease is the standard clinical example of a **subcortical dementia**, and its shape is examined precisely because it contrasts with the cortical pattern of Alzheimer disease. Cognitive decline broadly parallels the movement abnormalities and, like the earliest psychiatric features, it begins well before the disease is clinically manifest, emerging as much as 10 years before motor onset. The first casualties are cognitive flexibility and processing speed, so the pre-manifest patient struggles with task-switching and multitasking while more elementary functions are still intact.

A useful four-domain model orders the deficits by prominence in early to middle disease, and it is the ordering rather than any single item that repays memory.

MNEMONIC

Visuospatial, Executive, Memory, Verbal

The relative prominence of cognitive deficits in early to middle Huntington disease runs in that order: visuospatial deficits are the most pronounced, followed by executive dysfunction, then the less prominent memory and verbal deficits. Reading the domains in descending order captures why the picture looks nothing like an amnestic Alzheimer presentation.

Because the pathology is subcortical and frontostriatal rather than temporoparietal, the cortical hallmarks are muted: aphasia and agnosia are far less evident than in Alzheimer disease, and the patient relatively preserves language and recognition while losing speed, sequencing and

visuospatial competence. Tracked with a bedside instrument, the decline is gradual, with the MMSE score falling by roughly **1 to 2 points per year**. The MMSE and the MoCA are named here only as the tools used to follow the trajectory; their construction and scoring belong to the Psychotherapies, clinical assessment, MSE and nosology notes.

TRAP **Concluding that cognition is intact because memory and language are preserved.** Candidates primed on the amnestic, aphasic profile of Alzheimer disease look for those deficits, fail to find them, and miss a subcortical dementia that is already advanced in executive speed and visuospatial processing. In Huntington disease the deficit falls hardest on how fast and how flexibly the patient thinks; naming and recognition are relatively better preserved, though free recall is not spared.

14.6 Treating the chorea: VMAT2 inhibition and its alternatives

The disabling involuntary movements are treated by depleting presynaptic dopamine, and the mainstay agents are the **VMAT2 inhibitors**. Tetrabenazine and reserpine both act at the vesicular monoamine transporter to cause presynaptic depletion of catecholamines, the difference being that tetrabenazine is reversible while reserpine is irreversible. For the **chorea** of Huntington disease, tetrabenazine is begun at a starting dose of **12.5 mg daily**, and most patients can be managed on 50 mg a day or less. It is extensively metabolised by CYP2D6, and for the chorea, CYP2D6 genotyping is recommended for any patient who needs more than 50 mg daily. Its adverse effects include hypotension, parkinsonism and, critically, depression with suicidality.

The deuterated derivative, deutetrabenazine (Austedo), is FDA-approved specifically for chorea associated with Huntington disease and offers a smoother pharmacokinetic profile. For chorea in Huntington disease it is dosed across a range of 6 to 48 mg/day, started at 6 mg once daily and titrated upward in 6 mg/day steps at weekly intervals, with any total of 12 mg/day or more given in two divided doses and a maximum of 48 mg/day. Its schedule for chorea begins lower than the schedule used for tardive dyskinesia, a distinction worth keeping straight because the two indications share the drug but not the starting dose. The broader phenomenology of chorea, and the wider hyperkinetic movement-disorder overlap that includes Sydenham chorea, are addressed separately in this chapter; the concern here is its pharmacological control.

AGENT	MECHANISM	DOSE FOR CHOREA IN HUNTINGTON DISEASE	PRINCIPAL CAUTIONS
Tetrabenazine	Reversible VMAT2 inhibitor; presynaptic catecholamine depletion	Start 12.5 mg daily; most managed on 50 mg a day or less	Hypotension, parkinsonism, depression with suicidality; CYP2D6 genotyping if more than 50 mg daily
Deutetrabenazine	Deuterated VMAT2 inhibitor	6 to 48 mg/day, initial 6 mg once daily, maximum 48 mg/day	Black box warning for depression and suicidality
High-potency typical antipsychotic (haloperidol, fluphenazine)	Dopamine receptor blockade	0.5 to 2.0 mg/day	May worsen voluntary movement or apathy, or cause akathisia or excess sedation

Where a VMAT2 inhibitor is unsuitable, the chorea may also respond to low doses of a high-potency typical antipsychotic such as haloperidol or fluphenazine, and to benzodiazepines, to the dopamine-depleting agent reserpine, or to amantadine, which acts not by depleting dopamine but as an NMDA antagonist. None of these is free: as the next section sets out, the dopamine-depleting strategy carries its own affective cost.

14.7 The depression and suicidality loop in chorea treatment

The reason chorea management is a safety problem is that it collides with the suicide risk already established. Deutetrabenazine carries a **black box warning** for depression and suicidality in patients with Huntington disease, and tetrabenazine likewise can cause depression with suicidality. The clinical logic is uncomfortable but clear: the drug most effective against the movement disorder can worsen the very mood disturbance that already drives the elevated suicide rate. Deutetrabenazine is accordingly contraindicated in several situations that the clinician must be able to recite.

- In a Huntington patient who is suicidal, or who has untreated or inadequately treated depression.
- In hepatic impairment.
- In a patient taking a monoamine oxidase inhibitor.
- In a patient taking reserpine, tetrabenazine or valbenazine.

The practical rule that emerges is to reserve the VMAT2 inhibitors as a first-line choice for chorea for patients who are free of depressive symptoms and suicidal behaviour, and to monitor everyone on these drugs, together with their caregivers, for the emergence or worsening of depression, suicidality or unusual behaviour. The alternatives do not escape the dilemma. The dopamine-depleting agents, chiefly reserpine, increase the risk of depression in their own right, and the high-potency antipsychotics used for chorea may worsen voluntary movement or apathy, or add akathisia and sedation, so no route to chorea control is affectively neutral.

PITFALL

Starting tetrabenazine or deutetrabenazine to quieten distressing chorea in a patient who is quietly depressed or has voiced suicidal thoughts. It is an intuitive move, because the chorea is what is visible in the room, but in a patient with an untreated or inadequately treated depression it is contraindicated and can precipitate exactly the outcome the clinician most fears. Treat and stabilise the mood first.

In Huntington disease the suicide risk is high from the prodrome onward, and the VMAT2 inhibitors that control the chorea can deepen depression and suicidality, so never start one in a patient who is actively suicidal or whose depression is untreated or inadequately treated.

14.8 Course, prognosis and juvenile-onset disease

Huntington disease is relentlessly progressive and ultimately fatal, and framing the psychiatric care within that trajectory is part of managing it honestly. Death usually occurs 15 to 20 years after disease onset, and it comes most often not from the movement disorder as such but from its downstream consequences: the swallowing difficulties of advanced disease, which cause suffocation or aspiration pneumonia, and general debilitation. This shape matters for psychiatry because it defines the horizon against which decisions about aggressive symptom control, capacity, advance planning and the family's own bereavement are made, and it is why the depression of the middle years cannot be treated as though recovery to a normal life expectancy were on offer.

A distinct presentation deserves separate recognition. **Juvenile-onset Huntington disease** does not look like the adult chorea. It typically presents with bradykinesia, rigidity and dystonia, sometimes with seizures or myoclonus, and with chorea that is minimal or absent, and it follows a more rapid and more severe course than the typical adult form. The repeat length that underlies this earlier, more aggressive onset belongs to the genetics of the disease taken up separately in this chapter, and no numeric threshold defining juvenile onset is asserted here. Recognising the akinetic-rigid juvenile picture as Huntington disease at all is the clinical point, because the absent chorea can send the unwary toward an entirely different differential.

15 · Sydenham chorea, Tourette syndrome, Wilson disease and the hyperkinetic overlap

Three of the movement disorders a psychiatrist must know are united less by a shared pathology than by a shared danger: each can wear the mask of a primary psychiatric illness, and each is treatable when it is recognised for what it is. This section teaches Sydenham chorea, the post-streptococcal chorea of childhood, and the movement-disorder and neuropsychiatric lens of Wilson disease, and it sets both inside a broader hyperkinetic overlap that links post-infectious tics and obsessive-compulsive symptoms to the archetypal chronic tic disorder. Marks cluster here because the psychiatric surface of each condition misleads: emotional lability and obsessions in a choreic child, or a first psychotic episode in a young adult who in fact has a copper-storage disease. The diagnostic criteria of Tourette syndrome belong to the Child behavioural, emotional and other disorders notes and are named here, not reproduced;

the full hepatic and metabolic account of Wilson disease belongs to the Endocrine and metabolic psychiatry notes, and this section keeps to the brain.

15.1 Sydenham chorea: the only neurologic face of rheumatic fever

Sydenham chorea is the neurologic expression of acute rheumatic fever, and the exam-defining fact is worth stating plainly at the outset: it is the only recognised neurologic manifestation of rheumatic fever. It appears in about **20%** of patients with rheumatic fever, characteristically weeks to months after the throat infection that set the illness in motion, so that by the time the chorea declares itself the preceding streptococcal sore throat may be long forgotten. That latency is a trap for the unwary, who look for an acute infection and find none.

The disorder belongs to a family of post-infectious, immune-mediated diseases that follow infection with group A beta-haemolytic streptococcus. It helps to hold the whole set together, because they share a post-streptococcal, immune-mediated basis and differ in the organ the immune response strikes:

- acute rheumatic fever, of which Sydenham chorea is the sole neurologic manifestation;
- poststreptococcal reactive arthritis, striking the joints;
- PANDAS, a proposed disorder linked to the basal-ganglia circuitry of tics and obsessions; and
- poststreptococcal glomerulonephritis, striking the kidney.

For a psychiatrist the practical message of that list is that a recent streptococcal infection can express itself neuropsychiatrically, and that the interval between throat and brain can be long enough to hide the connection unless it is actively sought.

15.2 Molecular mimicry, the basal-ganglia circuit and the choreic picture

The pathophysiology is autoimmune and turns on molecular mimicry. Antibodies raised against the streptococcus cross-react with epitopes on the basal ganglia, and in doing so they disrupt the **cortical-basal ganglia-thalamo-cortical (CBGTC) circuits** whose task is to sculpt movement, mood and cognition into coordinated output. Because those loops carry motor, behavioural and cognitive traffic together, the clinical syndrome is never purely motor: the lesion that produces the chorea produces the psychiatric symptoms discussed below. Structural neuroimaging fits the model, showing volumetric increases in the caudate, putamen and globus pallidus, the very nuclei the antibodies target.

Clinically, Sydenham chorea usually declares itself between the ages of 5 and 15 years, with a female predominance of about 2:1. The onset of the movements is characteristically abrupt. They are choreoathetoid, restless and flowing, and they are accompanied by a striking hypotonia that can, in severe cases, leave the child floppy and weak; tics may occur, and when the disorder is severe the speech becomes dysarthric, slurred and explosive. The whole picture is easily mistaken for clumsiness, fidgetiness or a functional disorder before its post-infectious basis is recognised, which is exactly why the abrupt onset and the antecedent infection matter so much to the history.

15.3 The neuropsychiatric burden and its link to persistence

For the psychiatrist this is the important half of the disorder, because the behavioural disturbance is not incidental. Psychiatric symptoms are present in up to **70%** of children with Sydenham chorea, spanning emotional lability, anxiety, depression, obsessive-compulsive symptoms and features of attention-deficit hyperactivity, with psychosis occurring occasionally. Obsessive-compulsive symptomatology deserves separate mention: reported OCD prevalence reaches up to 40%, and it is higher in children whose chorea persists beyond 2 years. That correlation is instructive, tying the durability of the movement disorder to the durability of the psychiatric one and pointing back to a shared basal-ganglia substrate.

The overlap with obsessive-compulsive disorder is not coincidental. OCD is itself understood as a disorder of the cortico-striatal loops, so Sydenham chorea offers an unusually clean natural experiment in which a discrete immune insult to the striatum produces obsessions and compulsions alongside the movement disorder. The clinical corollary is simple: a child who presents with an abrupt obsessional illness together with a movement disorder is telling a single story, not two, and the obsessional symptoms should be screened for deliberately rather than waited for.

15.4 Protecting the heart and treating the brain

Management of Sydenham chorea sits largely in the outpatient setting, shared between paediatrics, cardiology and psychiatry, and it has two distinct targets that must not be run together. The first is cardiac protection. Because the chorea marks a child who has had rheumatic fever, that child needs long-term **benzathine penicillin** prophylaxis to reduce the risk of rheumatic heart disease. This is secondary prophylaxis against future streptococcal infection, not treatment of the chorea, and it is the single most consequential decision in the whole encounter, because the threat to life is cardiac rather than neurologic.

The second target is the neuropsychiatric syndrome itself, and here management is multimodal. It draws on cognitive behavioural therapy and SSRIs for the obsessional and affective symptoms, and, in more severe or refractory disease, on immune-directed treatments, namely corticosteroids, intravenous immunoglobulin and plasmapheresis, for the underlying autoimmune process, while antibiotics such as penicillin, a cephalosporin or amoxicillin treat any active streptococcal infection. It is worth being explicit about phase, because the targets change over time. The chorea itself is frequently self-limiting and often needs no specific suppression at all; when the movements are genuinely disabling, sodium valproate or carbamazepine is commonly used first, with dopamine antagonists or a VMAT2 inhibitor reserved for selected cases, so the Huntington regimen is not simply carried across. The obsessional and affective symptoms are treated over the medium term as they would be in any child, and the prophylaxis is the long-term commitment that outlasts every other element of care. Where a management question rests on formal guidance, name the reasoning and defer to current rheumatic-fever guidance rather than a remembered figure, because the allowed evidence base here fixes the principle of long-term prophylaxis without fixing a dose or an interval.

RULE **Sydenham chorea is the only recognised neurologic manifestation of rheumatic fever, and finding it obliges long-term secondary penicillin prophylaxis.** The chorea itself is largely self-limiting and its psychiatric symptoms are treatable, but the disease that produced them threatens the heart valves for years. Naming the chorea and stopping there, without instituting prophylaxis, misses the point of the diagnosis.

IN INDIA

Where rheumatic fever remains common, new-onset chorea in a child should place Sydenham chorea high on the differential and prompt its investigation, while other urgent causes of chorea, including Wilson disease, are excluded in parallel. Once Sydenham chorea is confirmed, the clinical priority is not the movement disorder but securing long-term secondary penicillin prophylaxis to protect the developing heart. Keep the diagnosis high in the differential of any child presenting with an abrupt movement or obsessional disorder, and examine, or arrange examination of, the heart before the encounter ends.

15.5 The hyperkinetic overlap: PANDAS, PANS and the tic disorders

Sydenham chorea sits at one end of a broader hyperkinetic overlap in which post-infectious immune activity, tics and obsessive-compulsive symptoms run together. Its closest relative is **PANDAS**, the paediatric autoimmune neuropsychiatric disorders associated with streptococcus, a proposed and still contested construct, which presents with obsessive-compulsive disorder and/or tics rather than with chorea. Its cardinal features are a prepubertal onset, an abrupt onset over roughly **24 to 48 hours** followed by a relapsing-remitting course, and a temporal relationship to streptococcal infection. The abruptness is the clinical signature: an obsessional illness that switches on almost overnight in a prepubertal child should raise the question.

The discriminator from Sydenham chorea that examiners favour is the movement itself, and it is set out in the accompanying comparison of the two syndromes.

AXIS	SYDENHAM CHOREA	PANDAS
Defining phenotype	Choreoathetoid movements	Obsessive-compulsive symptoms and/or tics
Age at onset	Childhood	Prepubertal, required
Course	Often self-limiting; can persist	Relapsing-remitting

The archetypal chronic tic disorder against which these post-infectious pictures are measured is **Tourette syndrome**; its diagnostic criteria and phenomenology belong to the Child behavioural, emotional and other disorders notes and are cross-referred here, not reproduced. The wider research construct of PANS, the paediatric acute-onset neuropsychiatric syndrome, makes the boundaries explicit, its criteria expressly excluding cases better explained by Tourette disorder, Sydenham chorea or systemic lupus erythematosus. Those exclusions are worth remembering because they define the overlap by what it is not: PANS is a diagnosis of exclusion that presupposes you have already thought of, and ruled out, the very disorders taught in this section.

TRAP **Collapsing PANDAS and Sydenham chorea into a single diagnosis.** They share a streptococcal trigger and a basal-ganglia target, but the movements separate them: the choreiform movements of Sydenham chorea almost always cause functional impairment, whereas those in PANDAS cause none. A child with disabling chorea is not a PANDAS case, and a child with abrupt obsessions and no functional movement disorder is not Sydenham chorea.

15.6 Wilson disease: copper, the ATP7B gene and the corneal ring

The second treatable masquerader in this group is a disorder of a single metal. **Wilson disease**, historically named hepatolenticular degeneration, is an autosomal recessive disorder in which mutations in the copper-transporting gene **ATP7B** leave the body unable to handle copper, which accumulates in the liver, brain and other tissues and drives hepatic cirrhosis and neuronal degeneration. For the neuropsychiatrist the important target is the brain, where the basal ganglia appear to be particularly affected, which is why the disorder earns its place among the movement disorders. Onset is typically in the second or third decade, with a mean age of onset of 17 years, and it is uncommon, with a frequency of about 1/30,000.

The one physical sign every candidate is expected to know is the **Kayser-Fleischer rings**, deposits of copper in Descemet's membrane at the corneal periphery, near the limbus, that may be seen on direct inspection but are found more reliably with a slit-lamp examination. Their great limitation is that they are often absent early in the disease, so their presence supports the diagnosis while their absence never excludes it.

PITFALL

Excluding Wilson disease because the Kayser-Fleischer rings are not there. They are often absent early in the disease, particularly in patients presenting with liver disease rather than neurologic signs, so a normal slit-lamp examination reassures nobody. If the clinical suspicion is present, pursue the copper studies regardless of the eye.

15.7 The movement disorder and the psychiatric masquerade

Wilson disease announces itself through one of three doors, and which door a given patient walks through decides who sees them first.

40%	40%	up to 20%
FIRST PRESENT WITH HEPATIC ABNORMALITIES	FIRST PRESENT WITH NEUROLOGIC SIGNS	FIRST PRESENT WITH PSYCHIATRIC ABNORMALITIES

The neurologic syndrome, which correlates with the copper load in brain and cerebrospinal fluid, is a movement disorder in the fullest sense, encompassing dysdiadochokinesis, dysarthria, tremor, dystonia, rigidity, choreoathetosis, bradykinesia, masked facies and micrographia; seizures occur in about 6%. The coexistence of hyperkinetic and hypokinetic signs in the same patient, tremor and dystonia alongside bradykinesia and a masked face, is itself a clue that the lesion sits in the basal ganglia rather than in a single transmitter system. On imaging there is increased T2 signal in the basal ganglia, including the putamen, and the midbrain may show the

nonspecific panda sign, a tegmental hyperintensity with relative preservation of the red nuclei and substantia nigra.

The psychiatric profile is where the disease earns its place in a psychiatry curriculum, because it can precede or overshadow the neurologic signs and it broadly resembles that of Huntington disease. The commonest features are personality change, that is irritability, aggressiveness and emotional lability, and depression, which can deepen into major depression with suicidality. Around **10%** of patients present with psychosis. Tellingly, psychiatric symptom severity correlates more strongly with the neurologic deficits than with the state of the liver, which underscores that this is a brain disease presenting to psychiatry, not merely a hepatic complication spilling over.

TRAP **Reading a young adult's first psychotic episode as primary schizophrenia or bipolar disorder without excluding Wilson disease.** Because the presentation can be purely psychiatric, the copper story is missed for years and the illness is misdiagnosed as schizophrenia or bipolar disorder. In a young patient with new psychosis, especially with any movement abnormality, personality change or hepatic history, the diagnosis must be actively excluded before it is called primary.

15.8 Making the diagnosis and clearing the copper

The diagnostic work-up follows the copper. Serum **ceruloplasmin** is below normal in 95% of patients, which makes it the natural first screen, but it must be read with two caveats: it is also low physiologically in newborns and in about 20% of heterozygous carriers, so a low value supports the diagnosis without clinching it. A raised copper concentration in a 24-hour urine collection is generally regarded as diagnostic, and liver biopsy showing increased hepatic copper remains the gold standard for confirmation. The tests form a ladder from an easy screen to a definitive but invasive confirmation, and in practice the diagnosis is assembled from the pattern rather than from any single result.

Treatment rests on a single principle: remove the copper, or stop it entering. The two chelating strategies mobilise stored copper so that it can be excreted, while blocking absorption prevents further loading from the diet, and the choice between them turns on the stage of disease and on tolerability rather than on any single rule. The urgency is real. Untreated, the disease is fatal, with death occurring within about 5 years, and outcomes are best when treatment begins early, before copper does irreversible damage, which is the whole argument for thinking of the diagnosis early. Because the disorder is autosomal recessive, a confirmed case also reframes the family: asymptomatic siblings share the same inheritance risk and are the people most likely to benefit from being found before their copper does its damage. This is one of the few movement and psychiatric disorders in which a specific and effective treatment exists, and in which the price of a delayed diagnosis is measured in irreversible neurologic damage and death. The detailed hepatic and metabolic management belongs to the Endocrine and metabolic psychiatry notes.

MNEMONIC

Chelate or block: penicillamine, trientine, tetrathiomolybdate, zinc

The two strategies against copper in Wilson disease: **chelate** and remove it with D-penicillamine, trientine or tetrathiomolybdate, or **block** its absorption from the gut with zinc salts.

GOOD TO REMEMBER

Wilson disease is the treatable one. Keep the threshold low to send a serum ceruloplasmin, a 24-hour urinary copper and a slit-lamp examination in any young adult presenting with an unexplained movement disorder or a first psychiatric illness, because early treatment gives the best outcome and delay costs irreversible damage. The test that starts the whole cascade is cheap; the diagnosis it catches is otherwise lethal.

In any patient in the second or third decade with an unexplained movement disorder or a first psychiatric illness, actively exclude Wilson disease: it is the treatable copper-storage disease that is otherwise fatal, and it is the diagnosis in this group you cannot afford to miss.

Autoimmune and infective encephalitides

16 · Anti-NMDA receptor encephalitis: course, antibody, tumour association and treatment

A young woman is brought in floridly psychotic, agitated and terrified, and within a fortnight she is seizing while her pulse and blood pressure swing out of control. The reflex of a psychiatric ward is an antipsychotic and a working diagnosis of a primary psychotic illness. That reflex, here, can miss a treatable neurological disease that opened with an entirely psychiatric face. The disorder that has taught a generation of psychiatrists this lesson is **anti-NMDA receptor encephalitis**, and the principle it drives home is simple: new psychosis in a young person that arrives with, or is quickly overtaken by, neurological signs is an autoimmune encephalitis until it is excluded. The marks cluster on this topic for three reasons worth separating. It is common, it is treatable, and it is the cleanest worked example medicine has of a single antibody producing a recognisable psychiatric syndrome. What follows defines the disorder, works through the antibody and the mechanism that ties it to psychosis and amnesia, describes who it strikes and the staged course it runs, sets out the tumour it so often conceals, and closes on the investigations and the immunotherapy that can reverse it.

16.1 The commonest autoimmune encephalopathy, and why psychiatry meets it first

Anti-NMDA receptor encephalitis, usually abbreviated NMDARE, is now thought to be the most common cause of autoimmune encephalopathy. That single fact reorganises the differential of acute psychosis in a young patient, because it means an antibody-mediated brain disease is not an exotic zebra to be considered last but a leading and eminently treatable cause to be considered early. What makes it a psychiatric problem before it is a neurological one is the way it announces

itself. A typically young adult female patient presents first with mood symptoms, agitation, psychosis or other variable psychiatric manifestations, and she may spend her first days under psychiatric rather than neurological care. The behavioural change, the paranoia, the disorganisation and the affective instability are convincing enough that a primary psychiatric label is the natural first guess.

The clinical weight of getting this right is asymmetric, and that asymmetry is the whole reason the disorder is taught so hard. A misjudged primary psychotic illness is merely treated late; a missed autoimmune encephalitis progresses to seizures, a movement disorder and autonomic collapse while the treatable window narrows. Because a treatable encephalitis can hide behind a psychiatric presentation and because it responds to immunotherapy, a new psychosis that carries any neurological colour warrants the antibody workup early. The cost of a paired antibody panel is trivial against the cost of watching a reversible encephalitis declare itself over the following weeks.

Commonest	Young women	Ovarian teratoma	Treatable
CAUSE OF AUTOIMMUNE ENCEPHALOPATHY	THE TYPICAL PATIENT, THOUGH ANY AGE OR SEX CAN BE AFFECTED	THE COMMONEST ASSOCIATED TUMOUR	MOST IMPROVE WITH IMMUNOTHERAPY

RULE **New psychosis in a young person that is joined by neurological signs is an autoimmune encephalitis until it is excluded.** The disorder is now regarded as the most common cause of autoimmune encephalopathy, and, unlike a primary psychotic illness, it is often reversible with immunotherapy once the antibody is cleared. The cost of missing it is high and the cost of testing for it is low, which is why the reflex to investigate should be automatic.

16.2 The antibody and how it unsettles the mind

The disease is defined by its antibody, and the mechanism is unusually satisfying because it explains the clinical picture rather than merely labelling it. Patients develop antibodies against the **GluN1 (NR1) subunit** of the NMDA receptor. The antibody does not simply flag the receptor for destruction; it binds the extracellular domain and causes the receptors to cross-link and be internalised, drawing them off the synaptic surface. The result is a reversible loss of functioning NMDA receptors, a state of **NMDAR hypofunction**, rather than the death of the neuron itself. That distinction carries the clinical message: because the neurons survive and only their surface receptors are stripped away, restoring the receptor population can restore the patient, which is the molecular reason the disease is treatable.

NMDA receptors sit at the centre of synaptic strengthening, so their loss has predictable consequences. The hypofunction produces defects in synaptic plasticity, and the failure of synaptic potentiation in particular maps onto the working-memory impairment these patients show. Because NMDA-receptor signalling is not confined to one region but woven through cortical, limbic and subcortical circuits, a single molecular lesion can express itself at once as disordered thought, disturbed memory, abnormal movement and autonomic dysregulation, which is why the clinical syndrome is so much broader than the tidy phrase of a receptor antibody would suggest. The same receptor hypofunction is the state implicated in glutamatergic

models of psychosis, which is why the psychiatric presentation can be indistinguishable at onset from a primary psychotic illness: the disease behaves, in effect, as a naturally occurring pharmacological model of psychosis produced by an autoantibody. Holding the mechanism in mind also frames the prognosis honestly. Most patients improve with immunotherapies, but recovery is not universal, and some are left with residual symptoms or a relapsing course even after the acute illness resolves.

16.3 Who it strikes

The archetype is worth learning precisely, and so is its exception. The characteristic patient is a young adult woman, and the disorder is decidedly commoner in that group, which is also the group in whom the tumour association is strongest. But the archetype is a centre of gravity, not a boundary. The disease can present in either sex and has been documented across the lifespan, from young children to older adults. An examiner who hands you a boy, a man or an older patient with a subacute psychosis and a movement disorder is often testing exactly this point: that the demographic is a clue, not a rule, and that a presentation which fits the syndrome should not be dismissed because the patient is the wrong age or sex.

The practical consequence is that the diagnosis rests on the pattern of illness rather than on demographics. A stereotyped psychiatric onset followed by neurological deterioration should trigger the workup whoever the patient is, while the young-woman archetype simply raises the prior probability and sharpens the search for an underlying tumour. Anchoring too hard on the typical patient is one of the reliable ways the diagnosis is missed.

16.4 The staged clinical course

In the older framing, NMDA receptor encephalitis is a variety of **paraneoplastic limbic encephalitis**, and like other limbic encephalitides it runs a recognisable staged course. It typically opens with several days of progressively severe amnesia and behavioural disturbance, the memory failure and change in conduct that first bring the patient to attention. This prodrome is then overtaken by the phase that defines the disease: psychosis, seizures, involuntary movements and autonomic instability which eclipse the earlier amnestic and behavioural picture. The tempo is the point. This is not a static psychosis but an evolving encephalitis, and the arrival of seizures, a movement disorder or swings in heart rate, blood pressure and temperature over days to weeks is the signature that separates it from a primary psychiatric illness.

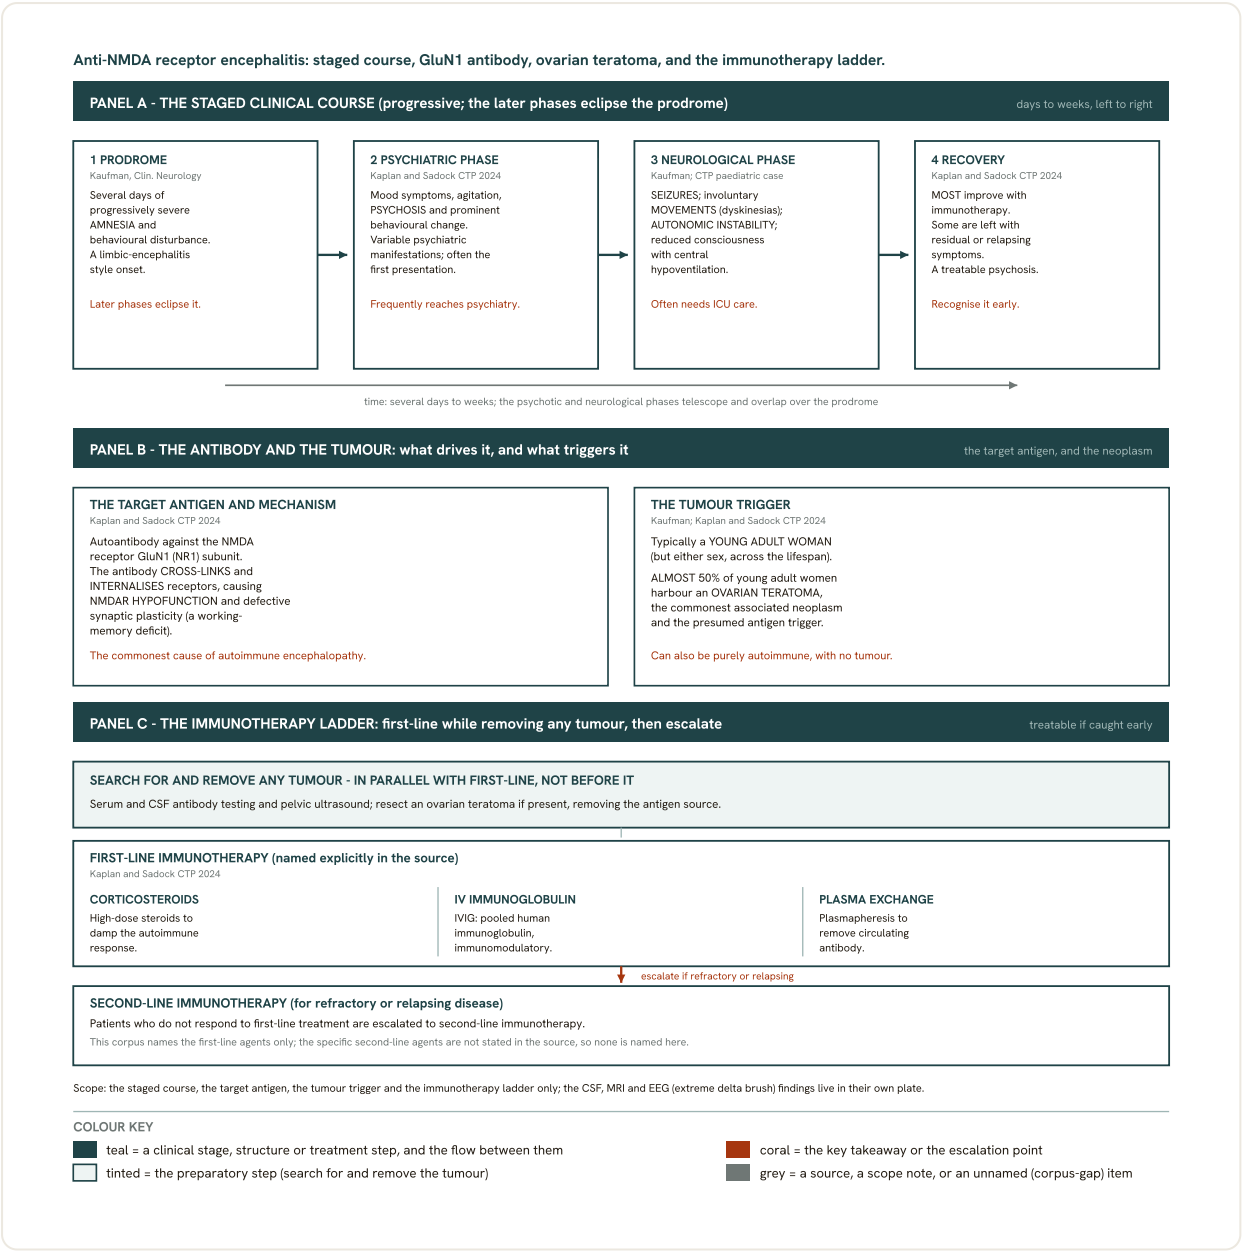


Fig 35.6 Anti-NMDA receptor encephalitis: the staged course and the treatment ladder. An autoantibody against the NMDA receptor GluN1 (NR1) subunit cross-links and internalises the receptor, producing NMDAR hypofunction; the illness runs a staged course from a prodrome of progressive amnesia and behavioural change, through a psychiatric or psychotic phase that often reaches psychiatry first, into a neurological phase of seizures, involuntary movements, autonomic instability and reduced consciousness that eclipses the prodrome, and then recovery, since most patients improve with immunotherapy. Almost 50 percent of young adult women harbour an ovarian teratoma, the presumed antigen trigger and the commonest associated neoplasm. Treatment starts first-line immunotherapy (corticosteroids, with intravenous immunoglobulin and/or plasma exchange) while searching for and removing any tumour, escalating to second-line immunotherapy in refractory or relapsing disease; the source names the first-line agents only, so no second-line agent is named here. (Kaufman, Clinical Neurology for Psychiatrists; Kaplan and Sadock CTP 2024.)

The autonomic phase is the dangerous one. Autonomic instability can escalate to the point of requiring intensive monitoring and support, and it is this stage, rather than the psychosis, that threatens life. The psychiatric phase itself frequently includes catatonia-like features; their formal phenomenology, the structured rating of catatonia and the distinction from primary catatonia belong to the Other psychotic disorders, delusional syndromes and catatonia notes and are named here rather than reproduced. The teaching point is the trajectory: a young patient who

moves from behavioural change and amnesia into psychosis, seizures, abnormal movements and dysautonomia over a short span is showing the natural history of this disease.

MNEMONIC

Amnesia, Psychosis, Seizures, Movements, Autonomic storm

The sequence in which the disease unfolds: an **A**mnestic and behavioural prodrome, then **P**sychosis, **S**eizures, involuntary **M**ovements and finally the **A**utonomic instability that eclipses the prodrome and endangers the patient.

16.5 The tumour it so often hides

The paraneoplastic label points to the search that must always follow the diagnosis. Overall, **almost 50%** of young adult women with the disorder harbour an **ovarian teratoma**, a germ-cell tumour that contains neural tissue and is presumed to trigger the antibody response by presenting nervous-system antigens to the immune system. The teratoma is the most common neoplasm associated with the disease, though a minority of patients are found to have other tumours, and, importantly, the syndrome can also occur as a pure autoimmune disorder with no underlying tumour at all. Finding and removing the teratoma is not incidental to treatment; it withdraws the antigenic drive and is part of definitive management.

The clinical corollary is that a search for the tumour is mandatory once the antibody is found, concentrated on the ovaries in the group most likely to carry one. The absence of a tumour on first imaging does not close the question, both because occult teratomas are easily missed and because the disease genuinely occurs without one, so the tumour search is sustained rather than performed once and forgotten.

TRAP **The tumour figure belongs to young adult women, not to every patient.** The frequently quoted teratoma association is the yield in young adult women specifically. Extend it to a man, a child or an older woman and it no longer holds, because it is the ovarian teratoma that anchors the number and those groups far less often carry one. A stem that hands you an older man and asks for the expected tumour yield is testing whether you kept the qualifier attached to the number.

16.6 Reading the investigations

The investigations are as much about not being falsely reassured as about confirming the diagnosis. The single most important habit is to distrust a normal scan. In most cases the MRI initially shows no abnormality, so an unremarkable early scan is exactly what the disease predicts and is never grounds to abandon the diagnosis. The cerebrospinal fluid is more informative. It usually shows a **lymphocytic pleocytosis**, the mark of an inflammatory process in the central nervous system, and it typically contains **oligoclonal bands**, which are nonspecific and appear equally in the CSF of multiple sclerosis and other inflammatory CNS illness; routine CSF can also be normal, and a bland tap does not exclude the disease.

The confirmatory test is the antibody itself, and it should be sent in both compartments. NMDA-receptor antibodies are found in both CSF and serum, but cerebrospinal fluid testing is the more

sensitive and specific of the two and carries the greater diagnostic weight, and CSF titres track the clinical course more closely than serum. Sending only one compartment is a mistake, because the two do not always agree, and a definite case can be positive in the CSF while negative in the serum. The electroencephalogram adds a suggestive but not decisive sign. The **extreme delta brush** pattern combines rhythmic or periodic delta slowing with overriding bursts of rhythmic beta activity at **20 to 30 Hz**, resembling the beta-delta complexes seen in neonates, with little reactivity or variability in sleep. It has been described in many patients with the disease, but it is explicitly nonspecific and its absence proves nothing.

INVESTIGATION	CHARACTERISTIC FINDING	WHAT IT DOES AND DOES NOT MEAN
MRI brain	Often normal at first	A normal early scan is expected and never excludes the disease
Routine CSF studies	Lymphocytic pleocytosis with oligoclonal bands	Supports CNS inflammation but is not specific, and normal routine CSF does not exclude the disease; oligoclonal bands also occur in multiple sclerosis
NMDAR antibody, paired serum and CSF	CSF IgG anti-GluN1 the more sensitive and specific; send paired serum too	Confirmatory; send both, but give the CSF result the greater diagnostic weight
EEG	Extreme delta brush	Suggestive but nonspecific; its absence does not exclude the disease

PITFALL

Treating a single antibody result as the end of the matter. Because cerebrospinal fluid testing is more sensitive and specific than serum, a definite case can be positive in the CSF while negative in the serum, so a diagnosis suspected on clinical grounds is not excluded by one negative compartment; a serum-positive but CSF-negative result should itself be read with caution. Sending serum and CSF together, and interpreting them together, is what avoids both traps.

16.7 Reversing it: immunotherapy and the tumour

Management follows directly from the mechanism, and it is genuinely a treatment of the disease rather than of the symptom. Care is inpatient, shared between neurology and psychiatry, and the sickest patients, those in the seizure and autonomic phase, need intensive monitoring and support. The target shifts with the phase of the illness: in the acute phase the priorities are controlling seizures and supporting autonomic function while the immune attack is brought under control; over the medium term the aim is to clear the pathogenic antibody and remove its source; and in the long term the concern is surveillance for relapse. The workup that opens management is serum and CSF antibody testing together with pelvic ultrasound, the imaging aimed squarely at the ovarian teratoma that may be driving the whole process.

The treatment itself combines removing the antigenic source with suppressing and clearing the antibody, and the components are worth listing because examiners ask for them by name.

- Removal of the tumour where one is found, which withdraws the antigenic drive and can itself turn the illness around.
- First-line immunosuppression with corticosteroids.
- Intravenous immunoglobulin, to neutralise and help clear the circulating antibody.
- Plasma exchange, to physically remove antibody from the circulation.
- Second-line immunotherapy for patients who do not respond to the first-line measures, the specific agents of which sit beyond the scope of these notes.

Steroids, with immunoglobulin and/or plasma exchange, are the first-line immunotherapy, deployed alongside tumour removal, and second-line immunotherapy is held for the non-responders. The reason this ladder works at all returns to the mechanism: the neurons are not dead, only stripped of surface receptors, so lowering the antibody load lets the receptor population, and the patient, recover.

GOOD TO REMEMBER

Once the diagnosis is suspected, do not wait for the picture to complete before acting. Send the paired serum and CSF antibodies and image the pelvis for a teratoma, and begin first-line immunotherapy while the tumour search proceeds; treatment and the hunt for the tumour run in parallel, not in sequence.

16.8 Prognosis and the psychiatric bottom line

The prognosis is the reason this disease rewards vigilance. Most patients improve with immunotherapies, which places NMDA receptor encephalitis among the small and precious group of psychoses that are genuinely reversible. But recovery is neither guaranteed nor always complete: some patients are left with residual deficits, and some run a relapsing course that requires continued surveillance and, at relapse, renewed treatment. Recovery can also be slow, unfolding over months rather than days, so the absence of an immediate response is not a reason to abandon immunotherapy. The corollary for counselling families is honest patience: improvement is the expectation rather than the exception, but it is measured in weeks and months, and an early plateau does not predict the eventual ceiling of recovery.

For the psychiatrist the lesson consolidates into a single habit of mind. The disorder is common, it wears a psychiatric mask, and it is treatable, so the differential of a young person with new psychosis and any neurological colour to the illness must include an antibody-mediated encephalitis, and the investigation must be sent rather than deferred. The whole clinical value of knowing this disease lies in acting on the suspicion early, because the treatment can only return what the antibody has not yet permanently taken.

Anti-NMDA receptor encephalitis is the commonest and a treatable autoimmune cause of new psychosis in a young patient: suspect it when a psychiatric onset is joined by neurological signs, confirm it with paired serum and CSF antibodies, hunt the ovarian teratoma, and treat with tumour removal and immunotherapy.

17 · The autoimmune limbic encephalitis spectrum

A patient in the fifth or sixth decade who over a few weeks loses recent memory, turns anxious or frankly psychotic and begins to have seizures has, until proven otherwise, an inflamed limbic system rather than a new primary psychiatric illness. The autoimmune limbic encephalitides are a family of immune-mediated inflammatory syndromes centred on the medial temporal structures, and they matter to the psychiatrist for a plain reason: their earliest and loudest features are psychiatric, so they arrive first in the psychiatry clinic and are readily mistaken for a mood, anxiety or psychotic disorder of middle-life onset. The marks concentrate here because the cost of that mistake is high. Many of these syndromes are driven by antibodies that are themselves the disease and are therefore reversible with immunotherapy, so a syndrome recognised is a syndrome that can be treated, while one dismissed as functional hardens into a fixed amnestic and cognitive deficit. This section sets out what defines the limbic syndrome and its mesial temporal substrate, the single organising distinction between the two classes of autoantibody, the specific antibody syndromes that carry named clinical signatures, and the tumour associations that make a new limbic encephalitis a reason to hunt for an occult cancer. Anti-NMDA receptor encephalitis, the most fully characterised member of the spectrum, and the red-flag approach to autoimmune psychosis are each addressed separately and are only named here.

17.1 The limbic syndrome and its mesial temporal substrate

The term **limbic encephalitis** describes an inflammatory process that settles on the limbic circuitry of the medial temporal lobes and expresses itself as a recognisable subacute syndrome: impairment of the formation of new memories, disturbance of mood and behaviour, psychiatric symptoms ranging from anxiety and irritability to hallucinations and delusions, and, very often, temporal lobe seizures. The tempo is the first diagnostic clue. Where a degenerative dementia evolves over years and a delirium over hours, the limbic encephalitides declare themselves subacutely, over days to weeks and sometimes across a couple of months, a march too fast for neurodegeneration and too sustained for delirium, and it is this cadence, more than any single symptom, that should prompt the question in the first place.

The anatomy explains the clinical picture and anchors the investigations. Whatever antibody or tumour underlies it, the disorder reflects medial, or mesial, temporal lobe dysfunction, so testing tracks that territory: magnetic resonance imaging may show acute swelling and signal change in the mesial temporal lobes, or atrophy later, the electroencephalogram may show spikes or slow waves emanating from the temporal lobes, and serum frequently carries antibodies reacting with limbic neurons such as anti-VGKC, anti-Hu or anti-Ma2 (Kaufman's Clinical Neurology for

Psychiatrists). Because the mesial temporal lobe houses the hippocampal machinery of episodic memory and sits at the heart of the limbic emotional circuit, the same lesion that erodes the ability to lay down new memories also destabilises affect and releases the temporal cortex into epileptic activity.

Mesial temporal	LGI1	CASPR2
MRI ATROPHY AND TEMPORAL-LOBE EEG DISCHARGES	AUTOIMMUNE CAUSE OF SIADH AND HYPONATRAEMIA	MORVAN SYNDROME AND AGRYPNIA EXCITATA

17.2 Two classes of antibody, and why the class predicts the outcome

The single most useful idea in this field is that the responsible autoantibodies fall into **two** mechanistically distinct classes, and that knowing which class you face predicts the tumour association, the reversibility and the response to treatment before any of those is formally established. The division is by the location of the target antigen. One class is directed against **intracellular** antigens, the classical onconeural targets such as anti-Hu, anti-Ma2 and anti-CRMP5. The other is directed against **extracellular, synaptic** antigens on the neuronal surface: the N-methyl-D-aspartate (NMDA) receptor, the AMPA receptor, gamma-aminobutyric acid (GABA) receptors, and the voltage-gated potassium channel complex antibodies, that is, LGI1 and Caspr2.

The mechanistic corollary is what makes the distinction worth memorising. An intracellular antigen is not accessible to a circulating antibody in a living neuron, so an antibody against it cannot itself be the effector of injury; it is a marker of a cytotoxic T-lymphocyte response that attacks the neuron directly and destroys it. Those syndromes are therefore tightly bound to an underlying cancer, tend to produce fixed neuronal loss, and respond poorly to antibody-directed immunotherapy because the damage is largely done by the time the diagnosis is made. A cell-surface synaptic antigen, by contrast, is exposed where antibody can reach it, so the antibody is itself pathogenic, typically by cross-linking and internalising the receptor or blocking its function. Because the neuron is disabled rather than killed, these syndromes are more often reversible, and they are the ones that reward immunotherapy. That is the whole reason the class matters at the bedside.

RULE **Classify the antibody by where its antigen sits, because the location predicts the biology.** Antibodies to intracellular antigens (anti-Hu, anti-Ma2, anti-CRMP5) are markers of a T-cell-mediated, cancer-associated, largely irreversible process; antibodies to cell-surface synaptic antigens are directly pathogenic and define the reversible, immunotherapy-responsive end of the spectrum. The target's address, not the antibody's name, is the organising fact.

FEATURE	INTRACELLULAR, ONCONEURAL ANTIBODIES	CELL-SURFACE, SYNAPTIC ANTIBODIES
Representative targets	Anti-Hu, anti-Ma2, anti-CRMP5	NMDA and AMPA receptors, GABA receptors, VGKC-complex (LGI1, Caspr2)
Role of the antibody	Marker of a cytotoxic T-cell response	Directly pathogenic effector
Antigen reachable in life	No, it lies inside the neuron	Yes, it sits on the cell surface
Typical reversibility	Usually fixed neuronal loss	Often reversible
Response to immunotherapy	Generally poor	Generally good

17.3 The VGKC-complex antibodies: LGI1 and Caspr2

Among the cell-surface antibodies, the group historically labelled anti-voltage-gated potassium channel antibodies deserves particular care, because the label is a partial misnomer that still appears on laboratory reports. The antibodies in this **VGKC-complex** do not, for the most part, bind the potassium channel pore itself; they bind proteins associated with the channel complex, and the two that matter clinically are **LGI1** and **Caspr2**. Recognising that a positive VGKC-complex result actually points to one of **two** named antibodies, each with its own clinical signature, is the difference between a vague immunological curiosity and an actionable diagnosis.

The two members separate along clinical lines. Caspr2 antibodies are the ones tied to the peripheral and autonomic overactivity syndrome, whereas LGI1 antibodies produce a predominantly limbic and amnestic picture and carry a characteristic metabolic marker of their own. The discipline that matters, for the examination and the clinic alike, is to resist collapsing the pair back into a single undifferentiated VGKC antibody once the specific target is known, because the target is what tells you which syndrome to expect and which associations to pursue.

> **PITFALL**
>
> Reading a positive VGKC-complex antibody as an antibody against the potassium channel itself. The clinically meaningful antibodies in the complex are directed at the associated proteins LGI1 and Caspr2, not the channel pore, and it is the specific target that carries the syndrome. Treating the older channel label as the diagnosis stops the workup one step short of the answer.

17.4 Anti-LGI1 antibodies and the clue of hyponatraemia

LGI1 antibodies have a metabolic fingerprint that is disproportionately useful to the psychiatrist, because it is cheap to find and easy to overlook. LGI1-antibody disease is one of the recognised autoimmune causes of the syndrome of inappropriate antidiuretic hormone secretion, and therefore of hyponatraemia (Localization in Clinical Neurology). In a patient presenting with a subacute amnestic and behavioural syndrome, an unexplained low serum sodium is a quiet pointer that the process may be autoimmune rather than primary psychiatric, and it is exactly the kind of routine result that is filed and forgotten once the working diagnosis has settled on depression or late-onset psychosis.

The pointer has to be read with the psychiatrist's own prescribing in mind, because the same source lists several drugs among the causes of inappropriate antidiuretic hormone secretion, and two of them are squarely psychiatric: carbamazepine and chlorpromazine both feature, alongside agents such as vincristine, chlorpropamide and cyclophosphamide. A hyponatraemia that emerges only after a mood stabiliser or an antipsychotic has been started is more likely iatrogenic than autoimmune, so the finding narrows the field only when the drug chart has been cleared first. Used that way, serum sodium becomes a small but genuine discriminator rather than a distraction.

GOOD TO REMEMBER

In a middle-aged patient with a subacute amnestic or psychotic presentation, check the serum sodium and read it in context. An autoimmune hyponatraemia driven by LGI1 antibodies is treatable, but the identical biochemistry follows carbamazepine or chlorpromazine, so clear the drug chart before you attribute a low sodium to the antibody.

17.5 Morvan syndrome and agrypnia excitata

The Caspr2 syndrome has a presentation that lands directly in psychiatric territory and is worth carrying by its eponym. When agitation combines with severe, intractable insomnia and signs of sympathetic overactivity, the constellation is called **agrypnia excitata**, and it should raise the question of **Morvan syndrome**, associated with antibodies to Caspr2 and, less commonly, LGI1. The name captures the essence: a state of pathological wakefulness driven from the brain rather than chosen by the patient, in which the normal architecture of sleep collapses and the patient is left agitated, sleepless and autonomically overdriven.

For the psychiatrist the hazard is the resemblance. A patient who has not slept for days, who is restless and agitated, and whose pulse, blood pressure and sweating betray an overactive sympathetic system looks, at first glance, exactly like severe agitated mania, an agitated depression or a stimulant state, and the reflex is to reach for sedation and an antipsychotic. The insomnia of agrypnia excitata, however, is not the reduced need for sleep of mania but an inability to generate sleep at all, and it keeps company with the autonomic and, often, peripheral neuromuscular features that betray its neurological origin. Reading the whole picture, rather than the mood-and-sleep fragment that happens to fit a primary diagnosis, is what turns the case toward the correct workup.

MNEMONIC

Agitation, Insomnia, Sympathetic overdrive

The three strands of **agrypnia excitata** in Morvan syndrome: **A**gitation, intractable **I**nsomnia and **S**ympathetic hyperactivity. Hold the triad together and the Caspr2 syndrome separates itself from the agitated states of primary psychiatry.

TRAP **Reading agrypnia excitata as agitated mania and treating it with sedation and antipsychotics alone.** The intractable insomnia, agitation and sympathetic overactivity of Morvan syndrome mimic a manic or mixed affective state, but the sleeplessness is a failure to generate sleep rather than a reduced need for it, and it travels with autonomic and peripheral nerve signs. Anchoring on the affective surface misses a treatable antibody-mediated disease.

17.6 The paraneoplastic associations and the search for a tumour

A limbic encephalitis can be the first announcement of a cancer that is otherwise silent, which is why a new subacute limbic syndrome is, among other things, a reason to look for a tumour. When the encephalitis is **paraneoplastic**, the antibody and the tumour tend to travel together, and the antibodies that accompany it include anti-Hu, anti-Ma, anti-amphiphysin, anti-VGKC and anti-GAD, alongside the anti-NMDA receptor antibody. The clinical value of the pairing is that the antibody hands you a shortlist of tumours to exclude, and treating a tumour that is found is essential and may stabilise the syndrome, though recovery is often limited in the intracellular, onconeural disease where cytotoxic T-cell damage is already done.

The tumours to exclude in paraneoplastic limbic encephalitis are a compact set worth committing to memory.

- Small-cell lung cancer.
- Testicular tumours.
- Breast cancer.
- Ovarian teratoma, the association that connects this spectrum to anti-NMDA receptor encephalitis.
- Thymoma, which also points toward the neuromuscular associations of the VGKC-complex antibodies.

Two of these deserve the extra beat. The ovarian teratoma link is the thread that ties the limbic spectrum to anti-NMDA receptor encephalitis, which is listed among the limbic-encephalitis antibodies and is addressed separately in its own right; the operational point to carry here is simply that a limbic syndrome in a young woman is a reason to image the ovaries. The thymoma link runs the other way, toward the peripheral and autonomic phenomena of the Caspr2 and VGKC-complex antibodies, and reminds you that the chest as well as the abdomen belongs in the tumour search.

17.7 Confirming the diagnosis: antibodies, imaging and the EEG

Confirmation rests on the convergence of three lines of evidence, each of which reflects the mesial temporal target of the disease. Structural imaging comes first. Magnetic resonance imaging may show atrophy of the mesial temporal lobes, and in the acute phase those structures may instead look swollen and signal-bright, so the scan is read for change in the limbic territory rather than for a single fixed appearance. The electroencephalogram is the second line, and it localises the irritable tissue, showing spikes or slow waves emanating from the temporal lobes and corroborating both the seizures and the anatomical focus. The third and most specific line is

serology: the serum frequently contains antibodies that react with limbic neurons, and a sensible panel samples both classes, including anti-VGKC, anti-Hu and anti-Ma2 among others.

Two habits of interpretation keep the workup honest. First, antibody testing is most informative when serum and cerebrospinal fluid are examined together rather than serum alone, because the compartment in which an antibody is found carries information the serum result cannot give by itself. Second, and more important at the bedside, these investigations are supportive and not perfectly sensitive early in the illness: an unremarkable scan or a non-specific tracing in the first days does not overturn a strong clinical suspicion, and this is a syndrome in which the tempo and the constellation of features carry as much weight as any single test. The diagnosis is assembled from the whole, not conceded on one normal result.

IN INDIA

Where the infective encephalitides dominate the differential of a subacute encephalopathy, the autoimmune limbic syndrome must still be actively tested for rather than assumed away, precisely because it is the treatable and often reversible member of the group. When herpes simplex, Japanese encephalitis, cerebral malaria and tuberculous meningitis have been addressed, an antibody-mediated encephalitis that was never considered is a reversible diagnosis missed. Keep it on the list, send the antibody panel, and let mesial temporal imaging and the electroencephalogram support the search rather than waiting for a textbook scan before believing it.

17.8 Why the psychiatrist is the first line of recognition

These often present with psychiatric features first and neurological ones second. A patient whose subacute change is dominated by anxiety, depression, psychosis, personality change or a disorder of sleep will be brought to a psychiatrist, and the window in which immunotherapy can reverse the illness is open at exactly that first contact. The clinician who holds the pattern in mind, a subacute course, prominent psychiatric features, memory impairment, seizures, and perhaps a quiet hyponatraemia or an intractable insomnia, is the one placed to act while action still helps. That is the whole argument for teaching this spectrum to psychiatrists rather than leaving it to neurology.

Two adjacent problems complete the map without belonging to this section. The specific red-flag features that should tip an apparently primary psychiatric presentation toward investigation for an autoimmune cause, and the immunotherapy by which such an autoimmune psychosis is actually treated, are developed separately and are not reproduced here. So too is anti-NMDA receptor encephalitis, the prototype of the antibody-mediated encephalitides, whose course, antibody and tumour associations and treatment are addressed in their own right. What this section asks the reader to carry is narrower and prior to both: the recognition that a new psychiatric syndrome in the wrong demographic, on the wrong timescale and with the wrong companions may be an inflamed limbic system asking for an antibody test.

A subacute psychiatric syndrome with memory impairment and seizures, especially alongside an unexplained hyponatraemia or an intractable insomnia, is a limbic encephalitis until the workup excludes it, and a single normal antibody panel or scan does not overturn a strong clinical suspicion.

18 · Autoimmune psychosis: red flags, the psychiatric differential and immunotherapy

A few psychoses are immunological emergencies wearing a psychiatric mask. A small but consequential group of patients reach a psychiatrist with what looks like a first episode of schizophrenia, an agitated mania, or a psychosis that will not settle on any antipsychotic, and are in fact suffering an **autoimmune or paraneoplastic encephalopathy**. The reason this topic earns marks is not that the syndrome is common but that it is treatable and eminently missable. The presentation is psychiatric, the cause is neurological, and the clinician who meets the patient first is usually the psychiatrist. This section teaches the one decision that matters: recognising when a psychiatric picture is atypical enough to warrant a search for neuronal antibodies, and knowing what that search and its treatment involve. It teaches the anti-NMDA-receptor-encephalitis-versus-catatonia distinction and the discrimination from a primary psychosis, and cross-refers the phenomenology and nosology that other chapters own.

18.1 Why the psychiatrist owns this question

The clinical logic begins with an observation from Kaplan and Sadock's *Comprehensive Textbook of Psychiatry*: these disorders should be suspected when a patient shows an atypical presentation of psychiatric illness, and above all when there are comorbid neurologic signs or symptoms. That pairing, an unusual psychiatric syndrome sitting alongside a neurological abnormality, is the whole of the **red flags** concept. The word doing the work is "atypical". A presentation is atypical when it departs from the natural history the clinician expects: the age is wrong, the onset is too abrupt, there is no prodrome, cognition is disturbed out of proportion to the affective or psychotic content, or the syndrome refuses to answer to treatment that should work. The presenting picture is itself protean, because an antibody-mediated disturbance of cortical and limbic circuits can surface as any of the functions those circuits carry: cognitive impairment, mood symptoms, frank psychosis, seizures, and movement abnormalities all appear, in varying combinations, and occasionally the illness presents with psychiatric symptoms and no neurologic symptoms at all.

The psychiatrist owns the recognition step because the referral pathway hands these patients to psychiatry before neurology ever sees them. What raises the stakes is that the underlying process is one most patients recover from once it is treated. The asymmetry of the two possible errors is what should set the threshold to investigate. Over-investigating a primary psychosis costs a venesection, a lumbar puncture, and a scan; under-investigating an encephalitis lets a treatable illness run as though it were incurable. When the downside of the two mistakes is that unequal, the rational clinician sets a low bar for ordering the antibody panel and accepts a high proportion of negative results as the price of not missing the positive one. The task, then, is not to make an

antibody diagnosis at the bedside but to hold the suspicion long enough to order the right investigations before a primary psychiatric label closes the case.

RULE **New or refractory psychosis carrying any neurological sign is an autoimmune encephalopathy until serum and CSF antibodies say otherwise.** The default in a young person with an abrupt, treatment-resistant psychosis and a soft neurological finding is not "unusual schizophrenia" but "encephalopathy not yet excluded". The exclusion is a laboratory and imaging exercise, not a clinical impression, because the treatable cause and the primary psychiatric one can look identical on interview alone.

18.2 The red-flag panel: what should stop you

Not every unusual psychosis needs an antibody screen, so the panel of features that should stop a clinician deserves to be memorised. The signal is tempo and neurology. A psychosis that evolves over days rather than months, that clouds cognition as it advances, or that brings seizures, abnormal movements, or autonomic disturbance is behaving unlike a primary psychiatric illness. The accompanying figure sets out this red-flag panel: the rapid-onset or treatment-refractory psychosis, the prominent neurological accompaniment, and the demographic and temporal cues that together justify a search for neuronal antibodies.

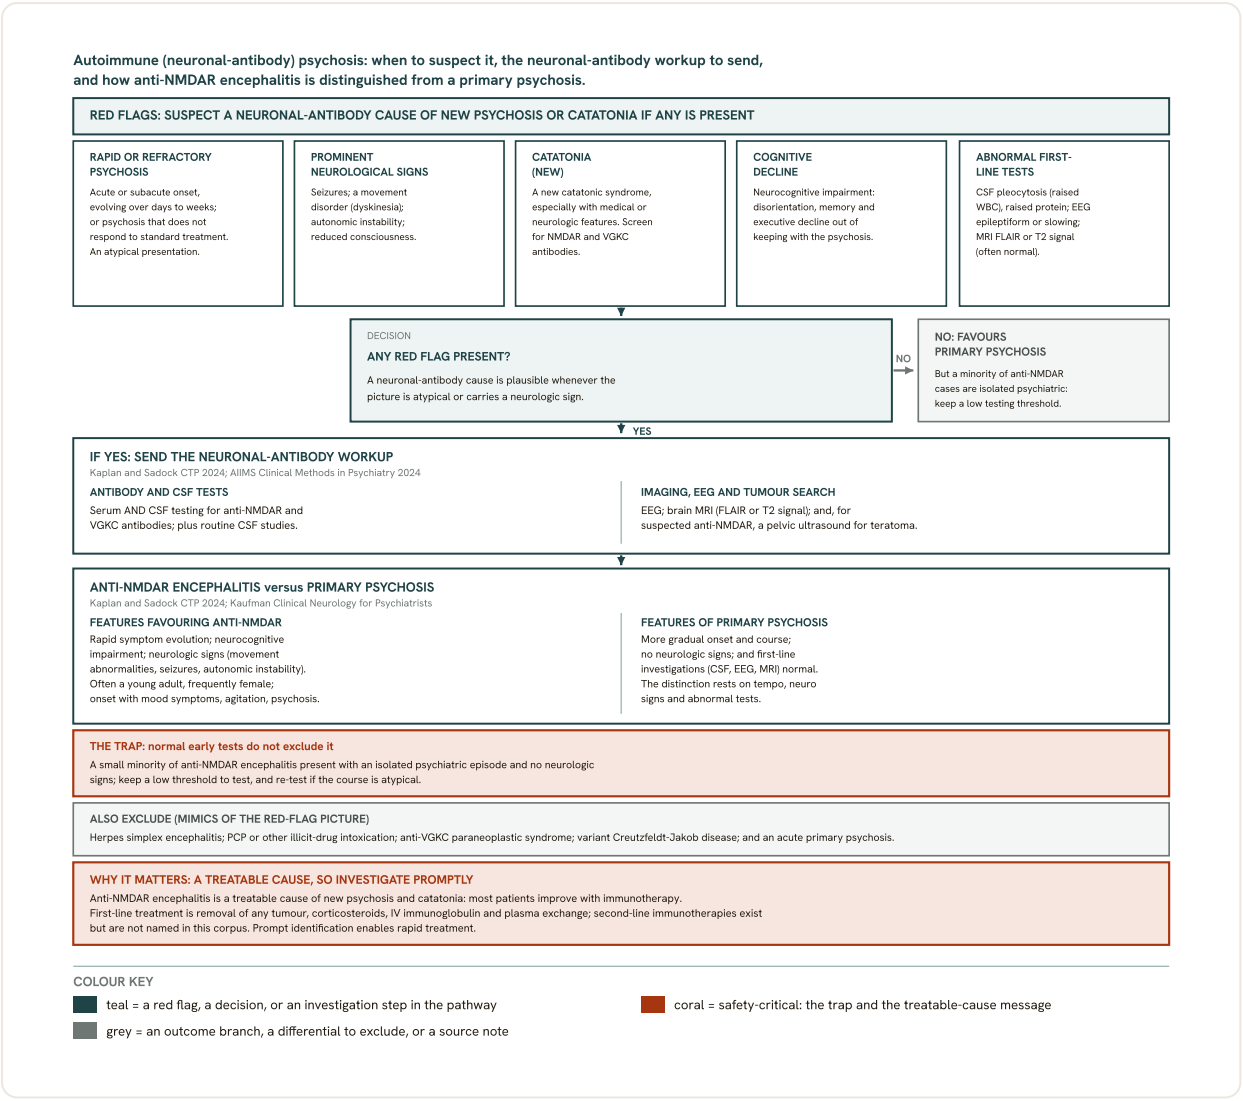


Fig 35.7 Autoimmune (neuronal-antibody) psychosis: the red-flag panel that prompts a neuronal-antibody search. New psychosis or catatonia that is rapid or treatment-refractory, carries neurologic signs (seizures, a movement disorder, autonomic instability, reduced consciousness), shows cognitive decline, or has abnormal first-line tests (CSF pleocytosis, an epileptiform or slow EEG, FLAIR or T2 MRI change) should trigger serum and CSF anti-NMDAR and VGKC antibody testing, routine CSF, EEG, MRI and a search for an underlying tumour. Anti-NMDAR encephalitis is told from a primary psychosis by its rapid tempo, neurocognitive impairment and neurologic signs, but a minority present with isolated psychiatric symptoms, so a treatable diagnosis is not excluded by normal early tests. (Kaplan and Sadock 2024; Kaufman's Clinical Neurology for Psychiatrists; AIIMS Clinical Methods in Psychiatry 2024.)

Read the panel as a trigger, not a diagnosis. Any one of its elements appearing in a patient who would otherwise be filed as a first psychotic episode is enough to move to investigation, because these features are not what a mood or psychotic disorder ordinarily generates and their presence, the neurological ones most of all, raises the prior probability. The individual cues cluster into a memorable spine.

MNEMONIC

Fast onset, Loss of cognition, Abnormal movements and autonomic signs, Generalised seizures

FLAG the psychosis that raises the red flag: **F**ast onset, the **rapid symptom evolution** that a primary psychosis rarely shows; **L**oss of cognition, an early neurocognitive disturbance out of keeping with the affective or psychotic content; **A**bnormal movements and autonomic instability; and **G**eneralised or focal seizures. Each of these is a discriminator that points the assessment away from a primary psychiatric diagnosis and toward an encephalopathic one, and although their weight varies, the neurological and paraclinical ones carrying the most, any single member is enough to keep the autoimmune question open.

18.3 From suspicion to the investigation set

Once an autoimmune or paraneoplastic encephalopathy is suspected, the investigation set is fixed and worth carrying as a unit. The *Comprehensive Textbook of Psychiatry* names four core studies to consider: serum and CSF antibody testing, routine CSF examination, the electroencephalogram, and brain MRI. Serum and CSF are both sampled because the two compartments do not always agree, and the antibody may be demonstrable in one when it is absent or equivocal in the other; a serum result alone can therefore mislead in either direction. The reasoning behind each study is what makes the panel memorable, because the yield is characteristic rather than dramatic and a clinician who expects a florid abnormality will be falsely reassured by a subtle one.

INVESTIGATION	WHAT IT MAY SHOW IN AUTOIMMUNE ENCEPHALOPATHY
Serum and CSF antibody testing	The neuronal antibody itself; sampled in both compartments because they may not agree
Routine CSF studies	Mildly to moderately elevated white cell count and protein, a modest inflammatory signal
EEG	Seizure activity, epileptiform discharges, or focal or generalised slowing
Brain MRI	Often normal; may show signal hyperintensities on FLAIR or T2-weighted sequences

The clinical grammar of the table is that a negative-looking result rarely closes the question. A modestly raised cell count is still an inflammatory finding, generalised slowing on the EEG is still an encephalopathic sign, and a scan reported as normal simply throws the diagnostic weight back onto the antibody and the CSF.

PITFALL

A normal brain MRI does not exclude autoimmune encephalitis. The *Comprehensive Textbook of Psychiatry* is explicit that imaging is often normal, so a clean scan reassures no one; the weight then falls on the clinical picture with the antibody and CSF results, not on the film. The same caution applies to a bland EEG. Treating a normal MRI as a negative test is how a treatable encephalopathy gets relabelled as a primary psychosis and discharged.

18.4 The autoantibody screen and the Indian workup

Indian teaching places the same investigations inside the everyday workup of a striking psychiatric presentation, which is what makes the screen practical rather than exotic. The AIIMS text *Clinical Methods in Psychiatry* sets out the laboratory investigations for catatonia and, among its special tests, lists CSF examination together with an **autoantibody screen** for N-methyl-D-aspartate (NMDA) receptor antibodies and voltage-gated potassium channel antibodies, alongside creatine phosphokinase, serum iron, brain imaging, an EEG to rule out epilepsy, and a drug screen. Sitting the antibody screen inside a standing investigation list, rather than leaving it to be remembered case by case, is precisely what stops it being forgotten in the patient who needs it most. The value is that the autoimmune question is asked routinely of a presentation the psychiatrist meets often, so the treatable cause is looked for before the psychiatric label has hardened.

IN INDIA

When a patient presents with catatonia or an atypical, treatment-refractory psychosis, keep the autoantibody screen for NMDA receptor and voltage-gated potassium channel antibodies, together with a CSF examination, on the investigation list rather than reserving it for the obvious neurological case. Following the AIIMS *Clinical Methods in Psychiatry* workup means the treatable encephalopathy is looked for as a matter of routine and not rescued late in the admission.

18.5 Anti-NMDA receptor encephalitis as an organic cause of catatonia

The pivot of this whole topic is that catatonia, so readily assumed to be psychiatric, has an organic differential the psychiatrist must actively hold. The *Maudsley Prescribing Guidelines* note that a catatonic syndrome is most commonly associated with psychiatric conditions such as bipolar disorder, schizophrenia, and major depressive disorder, but may also arise from a range of underlying medical conditions. Among those medical causes it lists autoimmune encephalitis, of which **anti-NMDA receptor encephalitis** is the archetype; systemic lupus erythematosus; and infections, especially of the central nervous system. The phenomenology of catatonia itself, the Bush-Francis rating, the lorazepam challenge, and malignant catatonia belong to the Other psychotic disorders, delusional syndromes and catatonia notes and are not re-taught here; what this section owns is the recognition that the antibody-mediated encephalitis sits on that organic list at all.

The practical consequence is that a catatonic presentation is never a licence to stop investigating. Reading catatonia as automatically confirming a mood or psychotic disorder is precisely how an antibody-mediated encephalitis is missed, because the psychiatric syndrome is genuine and the

organic cause hides behind it rather than replacing it. The medical differential is the reason the catatonia workup carries a CSF examination and an antibody screen in the first place, and it is why a patient who meets the criteria for catatonia still earns the full red-flag panel.

18.6 Telling autoimmune psychosis from a primary psychosis

When the encephalitis does declare itself psychiatrically, it does so in a recognisable way. The *Comprehensive Textbook of Psychiatry* describes a typically young adult female patient who initially presents with mood symptoms, agitation, psychosis, or other variable psychiatric manifestations, so at first contact the picture overlaps entirely with a primary psychiatric illness and the interview alone cannot separate them. The discriminators are not in the psychiatric symptoms but in what accompanies and surrounds them: a rapid evolution of symptoms, an early **neurocognitive impairment**, and neurologic signs such as movement abnormalities, seizures, and autonomic instability. It is the trajectory and the neurology, not the content of the delusions or the mood state, that break the tie. Two pulls toward the organic diagnosis are worth naming: the mismatch between how severe the syndrome is and how briefly it has taken to appear, and the company it keeps in a movement disorder or a seizure. Neither is proof, and each can occur in a primary disorder, but together they lift the prior enough to justify running the antibody workup in parallel with, rather than after, the steps taken to keep an agitated and confused patient safe. Waiting for the psychiatric picture to declare itself before investigating only delays the point at which appropriate treatment can begin. Where these accompaniments are present, the workup follows with serum and CSF antibody testing and with pelvic imaging in search of an underlying tumour, since an occult neoplasm can drive the process. The primary nosology of schizophrenia belongs to the Schizophrenia: management, antipsychotics and adverse effects notes, and that of depression to the Mood disorders: pharmacotherapy notes; the task retained here is only the discrimination.

> **GOOD TO REMEMBER**
>
> In a first psychotic episode that evolves over days, clouds cognition early, or brings a seizure or an abnormal movement, send serum and CSF for neuronal antibodies and, where anti-NMDA receptor encephalitis is suspected in a young woman, image the pelvis for a teratoma, before you commit to a primary psychiatric diagnosis. The demographic cue, a young woman with an abrupt psychiatric presentation, should raise rather than lower the index of suspicion.

18.7 The diagnostic trap and the mimics

The examination trap in this whole area is a single sentence. A small minority of patients present with an **isolated psychiatric presentation**, that is, psychiatric symptoms with no accompanying neurologic signs, so the reassuring absence of neurology does not exclude the diagnosis and the very feature a clinician leans on to feel safe is the one the syndrome exploits. Around this trap sits a differential the clinician must run before landing anywhere. Kaufman's *Clinical Neurology for Psychiatrists* notes that the picture may initially be mistaken for several other conditions, each of which can imitate the abrupt neuropsychiatric onset.

- Herpes simplex encephalitis, whose viral neuropsychiatry the chapter treats in its own right.

- Phencyclidine or other illicit drug intoxication, which can reproduce the psychosis and the dyskinesia.
- An anti-voltage-gated-potassium-channel antibody paraneoplastic syndrome, part of the limbic-encephalitis spectrum.
- Variant Creutzfeldt-Jakob disease.
- An acute, that is primary, psychosis.

TRAP **"No neurological signs, therefore not autoimmune."** Candidates anchor on the classic case with seizures and dyskinesias and forget that a small minority present with psychiatric symptoms alone. The examiner rewards the answer that keeps the antibody screen in play for the purely psychiatric presentation, because that is the patient the syndrome is most easily lost in.

A first psychotic episode that evolves rapidly, clouds cognition, or brings seizures or abnormal movements is an autoimmune encephalitis until serum and CSF antibodies say otherwise, and it is treatable.

18.8 The immunotherapy of autoimmune psychosis and its outcome

Management follows the same spine as any organic psychosis: where care happens, what it targets, and the modes available. The locus is inpatient and shared with neurology, because the definitive treatment is not psychiatric and the patient may need seizure and autonomic monitoring the psychiatric ward cannot provide alone. The focus is twofold, holding the acute psychiatric presentation safe while the disease-modifying treatment takes effect. The modes are dominated by **immunotherapy**. The *Comprehensive Textbook of Psychiatry* describes treatment as removal of the tumour if one is found, together with the **first-line** agents corticosteroids, intravenous immunoglobulin, and plasma exchange, and sometimes **second-line** immunotherapies. The corpus names the second-line only as a category and not by agent, so this section names it the same way and no further; the specific second-line drugs should be taken from current specialist guidance rather than assumed. Symptomatic psychiatric control, meanwhile, follows the principles set out in the Schizophrenia: management, antipsychotics and adverse effects notes, adjunctive to the immunotherapy and never a substitute for it. Read against the phases of care, the acute focus is the agitation, psychosis and any catatonic features that make the patient unsafe, held with the symptomatic measures the psychiatric team already commands; the mid-term focus is the immunotherapy course itself, delivered under neurology; and the longer-term focus is relapse surveillance rather than discharge on a first remission.

Serum and CSF	NMDA and VGKC	First-line	Most improve
ANTIBODY TESTING IS THE CORE INVESTIGATION	THE AUTOANTIBODY SCREEN IN A RED-FLAG CATATONIA	CORTICOSTEROIDS, IVIG, PLASMA EXCHANGE	ON IMMUNOTHERAPY; SOME RELAPSE

The outcome message is what makes the recognition worth the effort. The *Comprehensive Textbook of Psychiatry* states that most patients with the encephalitis improve with

immunotherapies, though some are left with residual or relapsing symptoms, and that the prompt identification of these autoimmune and paraneoplastic disorders facilitates the rapid initiation of appropriate treatment. That is the safe and defensible claim: the condition is treatable, and recognising it early lets treatment begin. How much a delay costs a given patient's recovery is a further question the source does not settle, and it should not be asserted beyond what the evidence supports. The teaching point stands on its own without that embellishment: find it, and most patients get better, which is exactly why the red-flag panel and the antibody screen are worth carrying in every atypical psychosis.

19 · Herpes simplex and the viral encephalitides with psychiatric sequelae

A survivor of encephalitis who is left amnesic, prone to seizures and altered in personality has almost always been injured in one place, the limbic system, and one virus reaches that territory more reliably than any other. **Herpes simplex encephalitis** is the reason a psychiatrist needs to know this disease well. It is the most frequent cause of serious, sporadic viral encephalitis, it has a mechanical preference for the parts of the brain that carry memory, emotion and social behaviour, and its survivors are handed to psychiatry with exactly the deficits that define the specialty. The marks here are not about culturing a virus. They reward a reader who can read an acute confusional illness for its temporal-lobe signature, explain why the damage produces the syndrome it does, and recognise a late autoimmune sequel that masquerades as relapse. The acute antimicrobial management of the infection itself belongs to neurology and infectious diseases and is directed by current guidance; what follows is the neuropsychiatry, which is where these patients live for the years afterwards.

HSV-1	Medial temporal	Haemorrhagic	Post-HSV NMDAR
THE COMMONEST SPORADIC, NONEPIDEMIC VIRAL ENCEPHALITIS	HIPPOCAMPUS, AMYGDALA AND UNCUS, DIENCEPHALON SPARED	INFERIOR TEMPORAL AND FRONTAL SURFACE, OFTEN BILATERAL	THE AUTOIMMUNE RELAPSE TO WATCH FOR

19.1 The virus, and the surfaces it attacks

The culprit is **herpes simplex virus type 1**, and the single most useful sentence a candidate can hold is that it is the most frequent cause of serious, sporadic viral encephalitis. The word that carries the weight is sporadic, or nonepidemic: unlike the seasonal and vector-borne encephalitides, this one arrives out of no outbreak and in no particular season, which is precisely why it must stay in the differential of any acute encephalopathy whatever the local epidemiology is doing. It is not a common disease in absolute terms, but among the individual, non-epidemic cases of viral brain infection it is the one you will meet.

Its behaviour inside the skull is stereotyped. The virus has a striking predilection for the undersurface of the frontal and temporal lobes, attacking the inferior and medial surfaces rather than the convexity. That anatomy is the whole clinical story in advance, because those surfaces carry the **limbic system**. The acute illness itself looks like most other encephalitides. There is fever, somnolence and delirium, the syndrome of an acutely inflamed brain, and the delirium here is no different in kind from the acute confusional states set out in the Stroke, TBI, delirium

and focal brain syndromes notes. What sets herpes apart is not the acute encephalopathy but where the inflammation settles, and therefore what it leaves behind. A febrile confusional illness that also brings new focal seizures, dense memory failure or a change of personality is not showing you a generic encephalitis; it is showing you the temporal lobe.

19.2 The neuropathology: acute necrotising damage to the medial temporal lobe

Pathologically the disease is an **acute necrotising encephalitis**, and both words matter. It does not simply inflame tissue and resolve; it kills it, which is why the deficits are so often permanent. The destruction is not diffuse but regional, showing a predilection for **the medial temporal lobe structures**. Postmortem and radiological studies converge on the same map: extensive lesions in the hippocampus, amygdala and uncus, the core of the medial temporal memory and emotion circuitry.

The negative half of that map is as examinable as the positive half. The diencephalon is left intact. This is the anatomical fact that separates the amnesia of herpes from the amnesias produced by deeper, midline injury, and it is worth fixing early, because a resident who knows only that both cause memory loss has not yet learned the discriminator. The virus works from the inferior surface inwards, sparing the thalamic and hypothalamic midline while it destroys the temporal lobe above it. Because the necrosis is haemorrhagic, the damaged tissue is not just non-functional but structurally disrupted, and that disruption is what the imaging and, in severe cases, the cerebrospinal fluid go on to reveal.

MNEMONIC

Hippocampus, Amygdala, Uncus

The three medial temporal structures that acute necrotising herpes simplex encephalitis destroys, while the diencephalon is spared. Fix the trio and you hold both the seat of the amnesia and the reason it is a temporal, not a diencephalic, memory loss.

19.3 The limbic signature: amnesia and complex partial seizures

Because the virus damages the limbic system, it routinely produces two neuropsychiatric consequences that follow directly from the anatomy. The first is **amnesia**. The medial temporal structures, and the hippocampus above all, are where new experience is turned into durable memory, so their destruction leaves the survivor unable to lay down new memories reliably, with older knowledge and general intellect relatively better spared. This is the memory loss of temporal-lobe injury, dominated at the bedside by failed new learning, although extensive injury can also erode remote and semantic memory.

Its counterpart is the amnesia of diencephalic damage, in which the lesion sits in the midline that herpes spares and the medial temporal lobe is intact. The two are anatomical mirror images, and holding that contrast is the point: herpes gives you a temporal, limbic amnesia precisely because the diencephalon is left untouched. The distinction is not academic. It tells you where the lesion is before the scan does, it predicts the shape of the deficit, and it stops a clinician from folding every dense amnesia into a single undifferentiated category. When a survivor of a temporal

encephalitis cannot retain new information, the anatomy has already told you why, and the confabulation, executive and gait features that cluster with midline injury are not what you should expect here. Charting the deficit formally is the work of the standardised cognitive instruments owned by the Psychotherapies, clinical assessment, MSE and nosology notes; the teaching here is only that the loss is temporal in origin and dominated by failed new learning.

The second limbic consequence is seizures. Damage to the temporal limbic cortex is epileptogenic, and it characteristically generates **complex partial seizures**, the focal, awareness-impairing seizures that arise from the mesial temporal lobe. In a survivor these are not incidental; the temporal focus that scarred with the infection remains an irritable one, and a chronic seizure liability is one of the commonest long-term legacies of the disease.

RULE **An acute encephalopathy with temporal-lobe localising features is herpes simplex encephalitis until proven otherwise.** Because herpes is the commonest sporadic viral encephalitis and it preferentially attacks the inferior and medial temporal surface, a febrile confusional illness accompanied by new complex partial seizures, dense anterograde memory failure, aphasia or a change of behaviour points to the temporal lobe, and therefore to this diagnosis, before any less common cause.

19.4 Kluver-Bucy features, aphasia and the frontal behavioural change

The behavioural sequelae are where herpes stops being a neurology problem and becomes a psychiatry one. Where the temporal destruction is one-sided the patient is left amnesic and seizure-prone, but when it involves **both temporal lobes** a further syndrome emerges. Bilateral temporal lobe damage in some survivors produces the human variety of the **Kluver-Bucy syndrome**, the behavioural picture of bilateral amygdala and anterior temporal injury, with altered emotional reactivity and changes in oral and sexual behaviour. The examinable discriminator is the laterality: this syndrome appears only when the damage is bilateral, and permanent temporal injury in this disease is often bilateral, which is why the syndrome is seen at all.

The inferior frontal involvement adds its own layer. Frontal and temporal lobe damage together can produce **aphasia** when the dominant hemisphere is affected, and, from the orbitofrontal and inferior frontal surface, the frontal lobe behavioural disorders of disinhibition, apathy and impaired social judgement. Put together, the survivor a psychiatrist meets may be amnesic, dysphasic, behaviourally disinhibited and emotionally altered, all from a single infection that respected the anatomy of the limbic and inferior frontal surfaces.

NEUROPSYCHIATRIC SEQUEL	ANATOMICAL SUBSTRATE	THE EXAMINABLE POINT
Amnesia	Medial temporal, hippocampus	Temporal, not diencephalic; failed new learning
Complex partial seizures	Mesial temporal irritative focus	A lasting seizure liability in survivors
Kluver-Bucy syndrome	Bilateral amygdala and anterior temporal lobe	Requires bilateral damage
Aphasia and behavioural change	Dominant temporal and inferior frontal surface	The personality change fielded by psychiatry

PITFALL

Reading the survivor as a primary psychiatric disorder. A patient left with the amnesia, disinhibition and altered emotional and sexual behaviour of temporal-lobe injury can be mistaken for a functional or psychotic presentation, especially when the acute illness was mild or was never recognised. The anatomical signature, a temporal amnesia with Kluver-Bucy features after bilateral injury, is what re-anchors the assessment on the brain rather than on the biography.

19.5 Investigation: imaging, the CSF and the electroencephalogram

The diagnostic effort is aimed at demonstrating the temporal-lobe process and excluding its mimics. In herpes and the other encephalitides the useful tests are a defined set:

- Lumbar puncture to examine the cerebrospinal fluid, which in the severe, haemorrhagic case may itself be revealing.
- Computed tomography, widely available and often the first scan to hand, though it is the less sensitive of the two structural tests early in the illness and magnetic resonance imaging is the preferred modality.
- Magnetic resonance imaging, which shows the temporal-lobe changes with far greater sensitivity.
- The electroencephalogram, useful where a temporal focus or a seizure tendency needs to be demonstrated.

On imaging the characteristic finding is hyperintensity in the inferior gyri of the temporal lobe, the radiological face of the same medial temporal disease. Herpes typically causes haemorrhagic inflammation of the inferior surface of the temporal and frontal lobes, and this is often bilateral, so the imaging mirrors the neuropathology structure for structure. In the most severe cases the temporal-lobe process becomes frankly destructive: haemorrhagic infarction allows blood to seep into the cerebrospinal fluid, so a bloody or xanthochromic tap in the right clinical setting is a corroborating sign rather than a red herring. The electroencephalogram contributes by localising, showing a disturbance over the affected temporal region that fits the rest of the picture, though a single recording never excludes the diagnosis. The logic of the investigation is worth making explicit, because each test answers a different question. Imaging asks where the tissue is damaged and shows the temporal-lobe process directly; the cerebrospinal fluid asks whether the illness is

inflammatory and, in the destructive case, betrays the haemorrhage; and the electroencephalogram asks whether that damaged temporal cortex has become electrically irritable. No single one of them is the whole answer, and the diagnosis is read from the pattern they make together against the clinical story rather than from any one result in isolation.

19.6 Managing the neuropsychiatric sequelae

Care runs across settings and phases. The acute illness is severe and is managed as an inpatient by neurology and infectious diseases; the antimicrobial treatment of the infection itself is directed by current guidance and is not reproduced here, and no regimen should be quoted from memory. The psychiatrist's work is what comes after, and it moves from the acute ward, through neurorehabilitation, to long-term outpatient follow-up.

The targets shift with the phase. In the acute and subacute period the problem is the encephalopathy and the emergent seizures. Over the longer term the picture is dominated by the amnesia, the chronic seizure liability, the behavioural change and the mood consequences of a life altered by brain injury. Across those targets the available modes are the familiar ones. The seizure disorder is managed as any structural focal epilepsy would be, under the relevant current guidance, which these notes do not restate. Symptomatic psychotropic treatment of the behavioural and affective sequelae is guided by the specific syndrome rather than by any protocol unique to herpes, and doses are not owned here. Cognitive rehabilitation, structured environmental and behavioural support, and sustained work with the family carry as much of the load as any drug, because a dense amnesia and a disinhibited survivor are managed at least as much by structure and support as by prescription. Where a guideline would ordinarily anchor a management section, the honest position is that the allowed evidence names none specific to this disease, so the reasoning is taught and the reader is pointed to current neurological and psychiatric guidance rather than to a remembered rule.

19.7 The late autoimmune sequel: anti-NMDA receptor encephalitis after herpes

The most important recent addition to this topic is a late complication that turns a recovering patient into a deteriorating one. Some patients, in the weeks after an apparent recovery from the viral illness, develop a fresh neurological and psychiatric decline, and case series have linked this to prior herpes simplex encephalitis. The entity is **anti-NMDA receptor encephalitis**, whose antibody, tumour associations and immunotherapy are set out among the autoimmune encephalitides covered elsewhere in these notes and are not repeated here. What belongs to herpes is the link and its clinical consequence.

The leading explanation is that the destruction of neural tissue by the virus exposes neuronal antigens to the immune system, which then mounts an antibody response against receptors on the surviving neurons; the mechanism is a hypothesis rather than a settled fact, but the clinical rule stands independently of it. A new movement disorder or a new psychiatric and behavioural deterioration arising after the acute infection has apparently settled should not be assumed to be recrudescent infection. It should prompt testing for neuronal surface antibodies, because the treatment that helps an antibody-mediated relapse is not the treatment of the virus. This is the

practical reason a herpes survivor who worsens weeks later is reviewed by the same clinicians rather than discharged as recovered.

TRAP **A fresh deterioration weeks after apparent recovery from herpes simplex encephalitis is not automatically a relapse of the virus.** Candidates and clinicians reflexively read a second decline as recrudescent infection and re-treat the virus, missing the post-herpes anti-NMDA receptor encephalitis that some case series describe. The tell is a new movement disorder or psychiatric change after recovery had begun; the answer is to test for the antibody, not to re-run the original assumption.

19.8 Holding herpes in the differential

The practical synthesis is a discipline about the differential. Herpes leaves a signature that is anatomical before it is serological: an acute confusional illness that carries temporal-lobe features, an amnesia of new learning, complex partial seizures, aphasia or a behavioural change, with imaging that shows the inferior and medial temporal surface. Recognising that signature is what earns the marks and, more importantly, what directs the investigation toward the temporal lobe and the cerebrospinal fluid.

GOOD TO REMEMBER

In a patient with an acute confusional illness and either new complex partial seizures or a dense failure of new memory, image the temporal lobes and examine the cerebrospinal fluid. The inferior temporal and frontal surface is where herpes leaves its mark, and in the severe, haemorrhagic case blood in the CSF is a corroborating clue rather than a distraction.

IN INDIA

The acute encephalitis syndrome here carries a broad endemic and epidemic differential that is set out elsewhere in these notes. The clinical discipline this section adds is to keep herpes simplex encephalitis firmly inside that differential whenever the picture is temporal, precisely because it is sporadic and nonepidemic: it does not announce itself with an outbreak or wait for a season. When a febrile encephalopathy brings temporal-lobe localising features, pursue the localising investigations, imaging of the temporal lobes and examination of the cerebrospinal fluid, rather than anchoring on the endemic causes alone.

Herpes simplex is the commonest sporadic viral encephalitis; it attacks the medial temporal and inferior frontal limbic surfaces, so survivors carry amnesia, complex partial seizures, Kluver-Bucy features and behavioural change, and a fresh deterioration weeks later is anti-NMDA receptor encephalitis until disproved.

Neurosyphilis and CNS infections

20 · Neurosyphilis and general paresis of the insane: stages, features and treatment

A treatable dementia hides inside this topic. Neurosyphilis is the neurological expression of an infection that medicine can cure outright, which is exactly why examiners keep it alive long after the disease itself became uncommon. It is the model of a reversible organic psychosis, and the candidate who recognises it changes the whole trajectory of the patient in front of them. It earns its marks because it sits at the junction of infection, dementia and florid psychiatric disturbance, because its physical signs are unmistakable once they are known, and because it remains a live differential wherever syphilis is still transmitted. This section follows the organism from its silent latency, through the early meningeal and vascular forms, to the two late parenchymatous syndromes of general paresis and tabes dorsalis, and closes on how the diagnosis is secured and treated.

20.1 The organism and the natural history

Neurosyphilis is caused by persistent infection with the spirochaete **Treponema pallidum**, and its neuropsychiatric syndromes are the delayed consequence of that organism seeding the meninges and the cerebral vessels early and, later, the substance of the brain and cord, stages that can overlap rather than following a fixed sequence. The epidemiology carries a strong HIV association: coinfection is actively looked for, and with advanced immunosuppression the disease can run a more aggressive course, though neurosyphilis is seen in HIV-positive and HIV-negative patients alike. A disease once associated with the untreated syphilis of the pre-antibiotic era is therefore looked for actively wherever syphilis is still transmitted, rather than waited for.

An asymptomatic stage, in which the organism is present in the nervous system but clinically silent, precedes overt disease by a variable latency. When neurosyphilis does declare itself, it does so in several symptomatic forms: acute syphilitic meningitis, meningovascular syphilis, general paresis and tabes dorsalis, each with its own tempo and its own signature, and these forms may overlap. The organising axis worth carrying is early versus late. Early neurosyphilis affects the coverings and blood vessels of the nervous system and appears in the first years after infection. Late neurosyphilis is **parenchymatous syphilis**: the organism has invaded the nervous tissue itself, and where it settles decides the syndrome, with infection of the brain producing general paresis and infection of the spinal cord producing tabes dorsalis. This anatomical logic is worth holding onto, because it predicts the clinical picture from first principles: a cerebral parenchymal infection presents as a dementia with psychiatric colouring, whereas a cord infection presents as a sensory and gait disorder, so the same organism produces two syndromes that look nothing like each other.

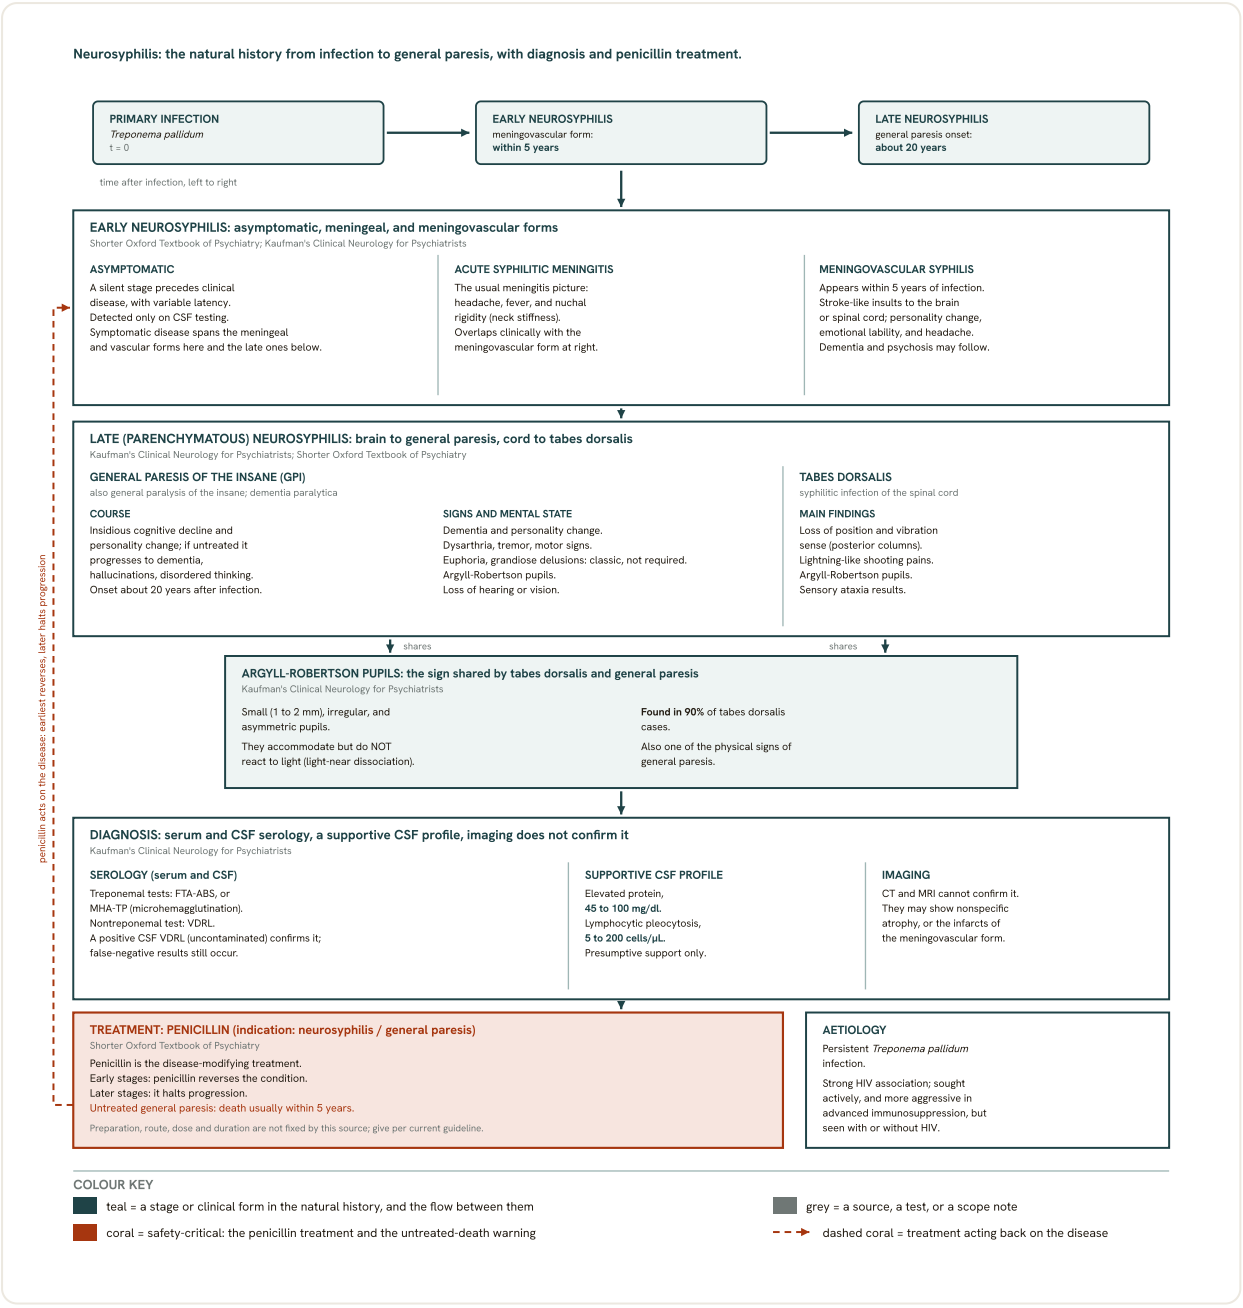


Fig 35.8 **Neurosyphilis across the natural history of syphilis.** Persistent Treponema pallidum infection runs from an asymptomatic stage through early neurosyphilis (acute syphilitic meningitis and meningovascular syphilis, the latter within 5 years) to late parenchymatous disease, in which brain involvement gives general paresis of the insane (dementia, personality change, dysarthria, euphoria with grandiose delusions, and Argyll-Robertson pupils, onset about 20 years after infection) and spinal-cord involvement gives tabes dorsalis; the small, irregular, light-near-dissociated Argyll-Robertson pupil is the shared sign. Diagnosis rests on treponemal and nontreponemal serology in serum and CSF with a supportive CSF profile, since imaging is nonspecific, and penicillin reverses early disease and halts later progression. (Shorter Oxford Textbook of Psychiatry; Kaufman's Clinical Neurology for Psychiatrists.)

The accompanying figure sets this natural history out from silent latency through the meningeal and meningovascular forms to the two late parenchymatous syndromes. The comparison below fixes the same map in a form the reader can scan at the bedside.

FORM	SITE OF INJURY	TYPICAL ONSET	HALLMARK
Acute syphilitic meningitis	Meninges	Early	Headache, fever, nuchal rigidity
Meningovascular syphilis	Cerebral and spinal arteries	Within 5 years	Stroke in the young; personality change, emotional lability
General paresis	Brain parenchyma	About 20 years	Dementia, personality change, dysarthria; florid affective and psychotic features
Tabes dorsalis	Spinal cord	Late	Sensory ataxia, lightning pains, Argyll-Robertson pupils

IN INDIA

Where syphilis remains in active transmission, a treatable spirochaetal cause belongs on the differential of any atypically young patient who presents with a dementia, an unexplained psychosis, or a stroke without vascular risk factors. The clinical decision is worth making a reflex: send a syphilis serology and examine the cerebrospinal fluid before the picture is written off as a primary psychiatric or neurodegenerative disorder, because this is the one such presentation that fully reverses when it is caught early.

20.2 Early neurosyphilis: meningitis and meningovascular disease

The earliest symptomatic form is **acute syphilitic meningitis**, which produces the usual symptoms and signs of meningitis - headache, fever and nuchal rigidity - and is easily missed simply because syphilis is rarely the first organism a clinician thinks of behind a lymphocytic meningitis. It shades into **meningovascular syphilis**, the form in which the inflammatory reaction centres on the walls of the cerebral and spinal arteries and sets up an obliterative endarteritis that narrows and occludes them. The two overlap in time and in mechanism, which is why they are considered together as the early, extra-parenchymatous face of the disease.

Meningovascular syphilis characteristically appears within 5 years of the primary infection, far sooner than the late parenchymatous forms, and it presents with two faces. It may declare itself abruptly as a stroke, an ischaemic insult to the brain or spinal cord occurring in a patient who is too young and too free of atherosclerotic risk factors for ordinary vascular disease. Alternatively it may creep in with personality change, emotional lability and headache. The stroke here is not embolic but the end result of that arteritic narrowing, an infarction in a vessel inflamed from without, which is why it can strike arteries and ages that atherosclerosis would ordinarily spare. Where the vascular injury accumulates, a dementia can follow, and that dementia may be coloured by psychotic symptoms, so the presentation reaches the psychiatrist as readily as the stroke physician. A stroke in the young without an obvious cause is one of the situations in which the syphilis serology earns its keep.

20.3 General paresis: latency, onset and course

The cerebral parenchymatous form, and the one that most concerns psychiatry, is **general paresis**, historically named general paralysis of the insane and also dementia paralytica. It is a disease of long incubation, starting to develop **about 20 years** after the original infection, so the patient who presents has usually long forgotten, or never knew of, the primary chancre. That latency is itself a diagnostic trap: nothing in the recent history points back to an infection acquired a generation earlier, and the connection has to be made from the clinical picture rather than the story.

The onset is insidious. The first changes are nonspecific cognitive impairments and a shift in personality that families register before clinicians do, and which are readily mistaken for depression, for an early dementia of another cause, or simply for ageing. If it is left untreated the process is relentless. The cognitive decline deepens into a frank dementia, and hallucinations and disordered thinking supervene. This drift from the subtle to the florid is what earned general paresis its old reputation as a great imitator of psychiatric illness, and it is the reason a subacute cognitive and personality change of uncertain cause should keep neurosyphilis somewhere on the list. The breadth of that differential is the point. General paresis can be taken in turn for a mood disorder when apathy and slowing dominate, for a primary psychosis when hallucinations and disordered thinking are prominent, and for a degenerative dementia when the cognitive decline leads; only the accompanying organic signs and the serology cut cleanly across all three.

20.4 The clinical picture of general paresis

Fully developed general paresis is a compound of a dementia, a personality change, dysarthria, and a scatter of motor symptoms and signs. Onto that organic base the illness may graft florid affective and psychotic phenomena. The classic portrait is of euphoria and expansive grandiose delusions, the patient of boundless wealth and importance, and it is this image that has fixed the disease in the examination memory. The mental changes seldom come alone, however, and it is the accompanying physical signs that betray the organic nature of the illness and steer the clinician away from a purely psychiatric reading.

TRAP **Grandiosity is the famous sign of general paresis, not a necessary one.** The Shorter Oxford Textbook of Psychiatry preserves the classic teaching of euphoria and grandiose delusions as the signature affective picture, and candidates duly memorise it as though its absence excluded the diagnosis. The examiner's counter is that Kaufman's Clinical Neurology for Psychiatrists records the opposite emphasis, that grandiose delusions, for all their notoriety, in fact arise only rarely. Treat grandiosity as a useful clue when it is present and never as a rule-out when it is not; the constant early thread is the nonspecific decline in cognition and personality, not the mood colouring laid over it.

The organic signature is a set of physical abnormalities that accompany the cognitive impairment: dysarthria, tremors, the Argyll-Robertson pupil, and loss of hearing or vision. The value of these signs is discriminative. A subacute dementia with personality change is a wide differential, but a subacute dementia with dysarthria, a tremor and an abnormal pupil is a narrow one, and it points hard toward an organic and specifically an infective cause. Examining for these

signs is therefore not a formality but the step that reframes the whole presentation. It is also the step that most often fails to happen, because a patient referred for grandiosity or a change in conduct is easily received as a psychiatric admission and never has the pupils, the speech and the gait examined at all.

20.5 Tabes dorsalis

The spinal parenchymatous form is **tabes dorsalis**, a disease of the dorsal columns and dorsal roots that produces a distinctive sensory syndrome rather than a dementia. Because the affected tracts carry proprioception and vibration, the patient loses the sense of where the limbs are in space, and the deficit expresses itself as a sensory ataxia: an unsteady, stamping, high-stepping gait that worsens sharply when the eyes are closed and the visual substitute for position sense is removed.

Its cardinal findings are:

- loss of position and vibration sensation, the deficit that underlies the sensory ataxia;
- lightning-like pains, brief lancinating shocks that shoot through the limbs or trunk;
- Argyll-Robertson pupils.

Tabes matters to the psychiatrist for two reasons beyond its own morbidity. It is the syndrome in which the Argyll-Robertson pupil is most reliably found, so recognising it anchors the pupillary sign, and it may coexist with general paresis in the same patient, so a dementia accompanied by prominent dorsal-column sensory loss should still raise neurosyphilis rather than being split into two unrelated diagnoses.

20.6 The Argyll-Robertson pupil

No physical sign is more firmly bound to neurosyphilis in the examiner's mind than the **Argyll-Robertson pupil**. The pupils are small, measuring 1-2 mm, and they are irregular and asymmetric between the two eyes. Their defining behaviour is a dissociation: they constrict briskly on accommodation, when the patient fixes on a closely held object, yet they fail to constrict to light. This **light-near dissociation** is the whole of the sign, and it is the origin of the bedside shorthand that the pupils accommodate but do not react.

MNEMONIC

Accommodates, but does not react

The Argyll-Robertson pupil keeps its near response and loses its light response: it **accommodates** to a close target **but does not react** to a light shone into it. Small, irregular and asymmetric complete the picture, and the four features together are far more specific than any one alone.

The sign is not merely characteristic but genuinely common in the spinal form. The constellation of small, irregular, unreactive pupils is found in **90%** of tabes dorsalis cases, which is why the pupils are the first thing examined when tabes is suspected. The same pupils appear in general paresis, so across neurosyphilis as a whole the Argyll-Robertson pupil is among the most useful physical signs available without any instrument at all.

20.7 Securing the diagnosis: serology, cerebrospinal fluid and imaging

The diagnosis rests on the laboratory, not the scanner. Serological testing is built on two classes of test that answer different questions. **Treponemal tests**, the fluorescent treponemal antibody absorption test (FTA-ABS) and the treponemal microhaemagglutination assay (MHA-TP), detect antibody directed against the organism itself and establish that infection has occurred. **Nontreponemal tests**, principally the VDRL, detect the reaginic antibody that tracks active disease. A VDRL that is positive in an uncontaminated cerebrospinal-fluid specimen confirms neurosyphilis.

PITFALL

A negative cerebrospinal fluid VDRL does not exclude neurosyphilis. The test is specific but insensitive in the fluid, and false-negative results are well recognised, so a strong clinical picture with a supportive fluid profile is treated and followed as neurosyphilis even when the VDRL there is negative. Reading a negative result as a rule-out is one of the commonest ways the diagnosis is missed.

The cerebrospinal fluid itself supports the diagnosis when it carries the marks of chronic meningeal inflammation: an elevated protein concentration, in the range of 45-100 mg/dl, and a lymphocytic pleocytosis of 5-200 cells/µL. These are supportive rather than confirmatory, and they belong to the presumptive picture that, taken together with the serology and the clinical syndrome, secures the diagnosis. Apart from a reactive cerebrospinal-fluid VDRL, no single result carries the diagnosis by itself, and the working discipline is to weigh the serology, the fluid and the clinical picture as one presumptive case rather than waiting for any one of them to be unequivocal before treating.

GOOD TO REMEMBER

Do not expect imaging to make the diagnosis. In neurosyphilis, CT and MRI cannot confirm the diagnosis: they may show nonspecific cerebral atrophy, perhaps weighted toward the frontal lobe, or the infarcts of the meningovascular form, but no pattern is specific to the infection. A scan is worth doing to exclude other causes, but a normal or nonspecifically atrophic scan neither confirms nor refutes neurosyphilis; the answer lies in the serology and the fluid, and time spent chasing an imaging signature is time lost.

20.8 Treatment and prognosis

The treatment is penicillin, and the single most important thing about it is how sharply the timing changes the result. Given in the early stages, penicillin reverses the condition; given later, once parenchymatous disease is established, it halts further progression but cannot restore what has already been lost. Care is led by the physician, while the psychiatrist manages the behavioural and psychotic disturbance in parallel across the acute and longer-term phases, but the disease-modifying step is antibiotic and it should not wait on the psychiatric picture settling first. The phases of care map cleanly: the acute task is to contain the behavioural disturbance and any agitated or psychotic features and to deliver the antibiotic; the medium-term task is to

consolidate cure and manage the residual cognitive and personality change; and the longer-term task is rehabilitative and social, supporting a patient whose late damage will not fully reverse.

RULE **Penicillin reverses early neurosyphilis and halts it later, so the entire clinical yield of the topic lies in remembering to test for it.** Because the organic damage of the late forms is only partly reversible, both the marks and the patient's outcome turn on making the diagnosis early. There is no version of this disease that is better left untreated, and none in which delay is neutral.

Untreated, the outlook is bleak. Once general paresis has declared itself, death usually follows within **5 years**. That figure is the counterweight to the reversibility above, the same disease resolving on early antibiotic yet killing within a few years when it is left alone. The precise regimen - the drug preparation, dose, route and duration - is a matter for current infectious-disease guidance rather than something this section can settle from its sources; what the examinable core fixes is the principle, that penicillin is curative early and protective late, not the milligrams.

~20 yrs	90%	5 yrs
LATENCY TO GENERAL PARESIS	ARGYLL-ROBERTSON PUPILS IN TABES	SURVIVAL IF UNTREATED

Any young patient with an unexplained dementia, psychosis or stroke earns a syphilis serology and a look at the pupils - neurosyphilis is the treatable one you cannot afford to miss.

21 · Neurocysticercosis and the endemic neuroinfections of India

The neuroinfections in this account are the ones an Indian psychiatrist genuinely meets, and each of them reaches the clinic wearing a psychiatric face. Neurocysticercosis, cerebral malaria and Japanese encephalitis together account for much of the seizure activity, cognitive decline, psychosis, delirium and encephalopathy that tropical central nervous system infection contributes to a psychiatric caseload. The marks gather here because the field rewards a clean map: pair each disease with the clinical clue that raises it, the investigation that confirms it, and the neuropsychiatric signature it leaves behind. The accompanying figure sets out that map before the detail is worked through.

Neuroinfections that reach the Indian psychiatrist: reading the clinical clue to the diagnosis.

Five endemic CNS infections, told apart by their clue, their neuropsychiatric picture, and the test that confirms them.

NEUROINFECTION	THE CLUE THAT REACHES YOU	NEUROPSYCHIATRIC PICTURE	HOW TO CONFIRM
NEURO-CYSTICERCOSIS (NCC) the calcified granuloma	Endemic across the Indian subcontinent. A parenchymal cyst at the grey-white junction, a few mm to 1-2 cm, then calcifies. About 30% of the epilepsy in endemic zones is attributed to it.	Seizures are the usual presentation and subside over years; psychosis usually complicates the epilepsy. A 4th-ventricle cyst gives Bruns syndrome: positional headache, vertigo, drop attacks.	CT screens, MRI characterises; serology for difficult cases. Over 50% relapse of seizures off medication; calcification predicts it. Antiseizure drug for seizures; antiparasitic and steroid if viable.
JAPANESE ENCEPHALITIS (JE) Asian encephalitis	A major encephalitis problem across Asia.	Encephalitic. Its impact is set beside measles encephalitis; survivors carry developmental and sensory sequelae.	Declared gap. This corpus gives no JE virology, vector, MRI signature or movement sequelae.
CEREBRAL MALARIA P. falciparum	Plasmodium falciparum. Retinal haemorrhages in 30-40% of cases.	An encephalopathy: impaired consciousness, delirium, or seizures.	Blood-smear microscopy; low CSF glucose favours it over viral encephalitis. Untreated usually fatal; treated mortality 15-20%. IV artesunate is first-line.
CNS TB the tuberculoma	Intracerebral tuberculomas, especially in the Indian subcontinent and in AIDS.	A mass lesion. On the subcontinent, one of the commonest causes of focal seizures, and may underlie new epilepsy in recent migrants.	CT or MRI can detect a tuberculoma even when the clinical picture does not suggest it. Declared gap: TB-meningitis syndrome and CSF profile.
SSPE measles, years later	A late measles complication, 7-10 years after infection; 85% present at a mean of 12 years, 15% as adults.	Behavioural change over weeks to years, then myoclonus (the hallmark), then dementia and akinetic mutism. Fatal: most die in 1-2 years, 95% within 5.	EEG: periodic sharp-wave complexes or burst-suppression. Raised CSF measles-antibody titre. Cowdry intranuclear inclusions on histology.

THE INDIAN-SUBCONTINENT READ (grounded in the body, not the health system) — Kaplan and Sadock 2024; Kaufman

Neurocysticercosis is endemic here, and about 30% of the epilepsy in endemic zones is attributed to it; and on the Indian subcontinent, tuberculomas are one of the commonest causes of focal seizures.

REFUSED FROM MEMORY, AND THE DECLARED GAPS — Corpus gap register: 213, 215, 216

The exam line that NCC is the single commonest cause of adult-onset seizures in India is not stated verbatim in a clean corpus source, so the grounded figures above stand in; JE virology and the tuberculous-meningitis syndrome are left as declared gaps.

COLOUR KEY

teal = the grounded clue, neuropsychiatric picture and confirming test for each infection

grey = a source note

coral = safety-critical: lines refused from memory and the declared corpus gaps

Fig 35.9 Endemic neuroinfections that reach the Indian psychiatrist, read from clinical clue to diagnosis. Neurocysticercosis (a calcified grey-white-junction granuloma, seizures the usual presentation, psychosis usually complicating the epilepsy), Japanese encephalitis (grounded here only as a major Asian encephalitis with measles-like outcomes), cerebral malaria (Plasmodium falciparum encephalopathy, retinal haemorrhages, blood-smear diagnosis, treated mortality of 15 to 20%), CNS tuberculoma (a mass lesion and, on the subcontinent, a leading cause of focal seizures), and subacute sclerosing panencephalitis (a late measles complication with myoclonus, periodic EEG complexes and a raised CSF measles titre). The exam cliche ranking neurocysticercosis first for adult-onset seizures in India is not stated verbatim in the clean corpus, so only the defensible figures are shown; Japanese-encephalitis virology and the tuberculous-meningitis syndrome are marked as declared gaps. (Kaplan and Sadock 2024; Kaufman's Clinical Neurology for Psychiatrists; Companion to Psychiatric Studies.)

Neurocysticercosis is the anchor of this section and is taught in full; cerebral malaria and Japanese encephalitis are the two other endemic neuroinfections taken up here. The chronic and subacute neuroinfections, including CNS tuberculosis, are addressed with the rare chronic infections elsewhere in these notes, and neurosyphilis has its own dedicated account; they are named here only where they belong in a differential.

21.1 The pork tapeworm and how its larva reaches the brain

Neurocysticercosis is disease of the central nervous system produced when the larvae of the pork tapeworm, **Taenia solium**, penetrate the CNS and encyst. Its geography is the geography of poverty and poor sanitation: it is endemic across most of the developing world, including the Indian subcontinent, most of Southeast Asia, sub-Saharan Africa, Central and South America and parts of China. For the clinician the decisive fact about transmission is how the larva reaches a human host. Infection is acquired by the faecal-oral route, through close contact with a Taenia solium carrier who is excreting eggs in the faeces, so the parasite passes from person to person

and the affected patient need not have eaten pork at all. A household or food-handler contact harbouring the adult tapeworm is the reservoir, which is why a contact history is often more revealing than a dietary one.

Once the eggs are ingested the embryos cross the gut wall, disseminate to the tissues, and the larval cysts reach their usual size in **2 to 3 months**. The illness itself, however, typically declares only years after the original infection, when the established cyst begins to degenerate and the host immune system turns on it. That long latency explains why the disease presents in adults with no recalled exposure and no acute prodrome, and why the parasite is so easily missed as the cause of a first seizure. The reason for the delay is immunological: a viable cyst is remarkably well tolerated, dampening and evading the host response so that it can sit quietly in the brain for years, and it is only when the larva dies, whether spontaneously or under treatment, that the surrounding tissue inflames and the disease becomes clinically loud. Holding on to that reversal, that the danger tracks the immune reaction to a dying parasite rather than the living one, is the single idea that runs through the whole of the presentation and the treatment.

21.2 The endemic burden: silent cysts and adult seizures

The public-health weight of neurocysticercosis is easy to underestimate because most infection is silent. In the endemic zones, seroepidemiological work finds specific antibodies in a substantial minority of the general population, and the healed footprint of old cysts, intraparenchymal brain calcification, is present on imaging in a considerable fraction of ordinary people scanned for other reasons; the great majority of these individuals never develop symptoms. What gives the disease its clinical importance is its contribution to epilepsy, and a large share of all epilepsy in the endemic regions is attributable to the parasite. The figures below fix the scale.

30% OF EPILEPSY IN ENDEMIC REGIONS ATTRIBUTABLE TO THE PARASITE	**> 10%** OF THE ENDEMIC POPULATION CARRY SPECIFIC SERUM ANTIBODIES	**10 to 20%** HAVE INTRAPARENCHYMAL CALCIFICATION ON CT, MOSTLY ASYMPTOMATIC

The mechanism that links the parasite to seizures is not the living cyst, which is largely tolerated, but its death: as the larva degenerates and calcifies it becomes an epileptogenic focus in the surrounding cortex. This is why a single calcified or degenerating granuloma is so often the substrate of an adult-onset seizure in an endemic setting, and why such a seizure deserves imaging rather than a reflex diagnosis of idiopathic epilepsy.

IN INDIA

Where cysticercosis is endemic, a first adult-onset seizure is imaged rather than assumed to be idiopathic epilepsy, because a solitary calcified or degenerating granuloma is a common substrate and its seizures are treatable. The clinical decision that follows from the biology is to obtain neuroimaging before committing a patient to a lifelong epilepsy label, and to hold the other great endemic mass lesion, the tuberculoma, in the same differential; that chronic infection is taken up with the subacute neuroinfections elsewhere in these notes.

21.3 Where the cysts lodge and the parenchymal natural history

The lesions of neurocysticercosis are small. Individual cysts range from a few millimetres to 1 to 2 centimetres in size, and they settle characteristically at the grey and white matter junction of the cerebral hemispheres, or within the deep grey matter. That distribution follows the blood supply, since the parasite is delivered haematogenously and lodges where perfusion is richest, and it is the reason the cortical and subcortical seizure focus is the commonest expression of the disease.

The natural history of purely parenchymal disease is comparatively benign. Patients whose lesions are confined to the brain parenchyma typically present with seizures only, and these tend to subside over the years as the cysts involute and calcify. A single cyst moves through a predictable arc, from a viable, fluid-filled vesicle that provokes little reaction, through a degenerating and inflamed phase in which seizures are most likely, to a small calcified scar that marks where it once sat; the calcification seen so often on imaging is simply the end point of that arc. This self-limiting parenchymal course stands in deliberate contrast to the extraparenchymal forms, in which cysts sit in the ventricles or the subarachnoid space, obstruct the flow of cerebrospinal fluid, and turn a nuisance disorder into a neurosurgical emergency. In the subarachnoid space the parasite can excite a dense inflammatory arachnoiditis that entraps vessels and cranial nerves, and in the ventricles it can block cerebrospinal fluid outflow to produce hydrocephalus, so the extraparenchymal disease threatens the patient by mechanisms the parenchymal form never reaches. It is these forms that carry the danger, and they are the subject of the next section.

21.4 The dangerous forms: cysticercotic encephalitis and Bruns syndrome

Two presentations move neurocysticercosis from an outpatient seizure problem to a life-threatening one. The first is **cysticercotic encephalitis**, in which a very large cyst burden provokes severe, diffuse cerebral inflammation with seizures and raised intracranial pressure. It occurs more often in children and young women, a distribution that reflects the intensity of the host immune response rather than any difference in exposure, and it is the immune reaction, not the parasite alone, that threatens the patient.

The second is **Bruns syndrome**, produced by a mobile cyst in the fourth ventricle. When the patient moves the head the cyst intermittently occludes the outflow of cerebrospinal fluid, a **ball-valve mechanism** that generates abrupt, positional surges of intracranial pressure. The clinical picture is characteristic and worth carrying whole:

- Episodic, severe headache.
- Vomiting.
- Vertigo.
- Sudden loss of consciousness precipitated by movement of the head.

TRAP **Reaching for an anthelmintic first in cysticercotic encephalitis.** When a heavy cyst burden is killed all at once the inflammatory storm intensifies, so in this fulminant form antiparasitic treatment given on its own can worsen the illness and prove fatal. The swelling and the host immune reaction are controlled first, and cyst-directed therapy is withheld or deferred until it is safe.

21.5 The neuropsychiatric presentations

The reason neurocysticercosis lands on a psychiatric desk is that its cortical footprint disturbs cognition and, through the epilepsy it causes, mental state. Cognitive impairment is common and spans a wide range, from a mild reduction in cognitive performance through to a severe dementia, depending on lesion burden, location and the seizure history. The formal characterisation of that impairment, with the standardised bedside cognitive instruments, belongs to the Psychotherapies, clinical assessment, MSE and nosology notes; what matters here is that a dementing picture in a patient from an endemic region should prompt a search for calcified cysts rather than an assumption of a primary neurodegenerative process. Because a share of the cognitive disturbance is driven by the seizures and by active inflammation rather than by fixed tissue loss, some of it is potentially modifiable, which is a further reason not to close the assessment at a diagnosis of irreversible dementia.

Psychotic episodes also occur, and the best current reading is that they arise largely as a complication of the epilepsy rather than as a direct parasitic psychosis. The distinction is practical: it directs attention to seizure control and to the peri-ictal and interictal mechanisms of psychosis rather than to the cyst itself. The primary nosology and management of a psychotic illness are set out in the Schizophrenia: management, antipsychotics and adverse effects notes, and of depressive illness in the Mood disorders: pharmacotherapy notes; the point in this chapter is only that this population reaches those diagnoses through an organic, epilepsy-driven route that must be recognised first.

21.6 Diagnosis and the principle of treating the inflammation

Diagnosis rests on neuroimaging. CT of the head is used to screen, MRI of the brain characterises the lesions more precisely, and serology is reserved to confirm the diagnostically difficult case. The imaging also stages the disease, because the number, the location and the calcified or active state of the cysts drive every subsequent decision. Management then runs on two tracks that must be coordinated. The seizures are managed where the patient can be seen safely, in most cases as an outpatient, and usually respond to monotherapy with a first-line antiseizure medication. The underlying infection and its inflammatory consequences are managed according to burden and form, with the fulminant and ventricular presentations of the previous section requiring inpatient and often neurosurgical care for cerebrospinal fluid diversion.

The shape of care follows directly from that staging. Uncomplicated parenchymal disease is managed in the outpatient setting, while the fulminant encephalitic and obstructing ventricular forms need inpatient, emergency and neurosurgical care. The targets shift with time: acute seizure control, and where intracranial pressure is raised the relief of that pressure, come first; eradication of the parasite and control of the inflammation follow; and the prevention of seizure

relapse governs the long term. Across those phases every available mode of treatment is in play, from drug therapy for the seizures, the inflammation and the parasite, through the procedural diversion or removal of an obstructing cyst, to the social and rehabilitative support a patient left with epilepsy or cognitive impairment will need.

RULE **In neurocysticercosis you are treating the host inflammatory response as much as the parasite.** Because it is the death of the cyst that inflames the brain, corticosteroids are used to damp the immune reaction while anthelmintics such as praziquantel and albendazole clear the infection, and antiparasitic therapy is timed and covered by steroid rather than given reflexively. Where a management decision rests on the number and viability of cysts, individualise it against current expert guidance rather than a fixed rule.

GOOD TO REMEMBER

Do not be quick to withdraw antiseizure medication. **More than 50%** of patients relapse after their treatment is stopped, and the risk is higher still with multiple lesions, a history of recurrent seizures, or brain calcification. These features argue for a longer, more cautious course rather than an early trial off medication.

21.7 Cerebral malaria: the falciparum encephalopathy

Of the malaria parasites, **Plasmodium falciparum** carries the particular risk of prominent neuropsychiatric involvement, the syndrome of **cerebral malaria**. It presents as an encephalopathy, with impaired consciousness, delirium or seizures, and it sits within a disease of enormous scale: malaria as a whole affects an estimated 240 million people each year and kills over 500,000, and one study found neurologic involvement in 48% of children with malaria, the risk concentrated in the endemic areas of sub-Saharan Africa. The assessment of the delirium itself is developed in the Stroke, TBI, delirium and focal brain syndromes notes; the task here is to recognise its infective cause.

A physical sign helps at the bedside: retinal haemorrhages are found in **30% to 40%** of cerebral malaria, so the fundus is part of the examination. Diagnosis is confirmed by blood smear microscopy, and where the cerebrospinal fluid is examined a low CSF glucose favours cerebral malaria over a viral encephalitis. Treatment is urgent, because the outcome hinges on it: untreated, cerebral malaria is generally fatal, whereas with treatment mortality falls to **15% to 20%**. Intravenous artesunate is the treatment of choice, alongside the neuropsychiatrically relevant management of cerebral oedema and seizures. In the treated survivor the encephalopathy is often reversible, which sharpens the imperative to reach the diagnosis quickly: this is an illness in which the disturbance of consciousness and behaviour can lift with prompt antimalarial and supportive treatment, and in which delay converts a survivable encephalopathy into a fatal one. For the psychiatrist the practical consequence is that the febrile, confused patient in an endemic setting is a medical emergency to be worked up, not a behavioural problem to be sedated.

PITFALL

Settling on a viral encephalitis, or a primary psychiatric disturbance, in a febrile encephalopathic patient without examining the fundus or reading the CSF glucose. Retinal haemorrhages are present in a large minority of cerebral malaria and point straight to the diagnosis, while a low CSF glucose favours malaria over a viral cause; missing both delays the one treatment that changes the outcome.

IN INDIA

In an endemic area, a febrile encephalopathy, an acute delirium, or a first seizure with fever is investigated with blood smear microscopy before it is attributed to a viral encephalitis or a primary psychiatric disturbance, because falciparum malaria is the one cause here that is both rapidly fatal and rapidly treatable. The clinical decision that follows is to send the smear early and to begin specific antimalarial treatment without waiting once the diagnosis is secure.

21.8 Japanese encephalitis and the endemic differential

The third of the endemic neuroinfections is **Japanese encephalitis**. Across Asia it is a major public-health problem, and the sources used here place its impact alongside that of measles encephalitis, a leading cause of learning disability, sensory impairment and under-five mortality in low-income countries. For the psychiatrist the relevance lies in the survivors, who carry a heavy burden of developmental and sensory sequelae after the acute illness. The detailed virology, vector and movement-disorder aftermath of Japanese encephalitis are not developed from the clean sources available for this account, and are deliberately not asserted here; the reader who needs them should turn to current references.

Holding the three endemic neuroinfections against one another is the most useful way to keep them straight, because each is separated from the others by its clinical clue, its confirmatory test and the psychiatric picture it produces.

ENDEMIC NEUROINFECTION	CLINICAL CLUE	CONFIRMATORY INVESTIGATION	NEUROPSYCHIATRIC SIGNATURE
Neurocysticercosis	Adult-onset seizures, calcified or cystic lesion	CT to screen, MRI to characterise, serology to confirm	Seizures, cognitive decline to dementia, epilepsy-related psychosis
Cerebral malaria	Febrile encephalopathy, falciparum infection, retinal haemorrhage	Blood smear microscopy; low CSF glucose	Impaired consciousness, delirium, seizures
Japanese encephalitis	Major viral encephalitis of Asia	Viral-encephalitis workup	Developmental and sensory sequelae in survivors

In endemic India a new adult seizure with a calcified brain lesion raises neurocysticercosis strongly, but is imaged and worked up against its mimics, the tuberculoma above all; a febrile encephalopathy is treated as cerebral malaria, with other encephalitic causes pursued alongside, until repeated blood smears are negative; and the treatment rules diverge, steroid-covered antiparasitic therapy for selected neurocysticercosis but urgent intravenous artesunate and supportive care, not routine steroids, for cerebral malaria.

22 · SSPE, CNS tuberculosis and the rare chronic neuroinfections

Some infections of the central nervous system announce themselves not with fever and a stiff neck but with a change in behaviour, a slipping intellect or a first psychotic episode, and they do so months or years after the organism first arrived. This section gathers the subacute and chronic neuroinfections that a psychiatrist meets in that disguise: subacute sclerosing panencephalitis, the late and lethal sequel of measles; the tuberculoma, an intracerebral mass that is a leading structural cause of seizures across the Indian subcontinent; and Whipple's disease, a systemic infection whose one pathognomonic sign rescues an otherwise missed and treatable dementia. Marks cluster here for two reasons. The first is that each condition wears a psychiatric or cognitive mask, so the diagnosis is made by the clinician who thinks of it rather than by the one who waits for it. The second is that several are treatable, which makes missing them costly in a way a purely degenerative diagnosis is not. Progressive multifocal leukoencephalopathy is named here only in passing, because its clinical account belongs with the CNS opportunistic infections of HIV; the syphilitic and endemic parasitic neuroinfections are likewise taught in full elsewhere in these notes and are set beside this group only to complete the differential.

22.1 Subacute sclerosing panencephalitis: persistent measles in the brain

Subacute sclerosing panencephalitis (SSPE) is a progressive and ultimately fatal panencephalitis that develops long after an apparently uncomplicated measles infection, and it is thought to arise from persistence of the virus within the central nervous system rather than from a fresh insult. The causative organism is measles, or rubeola, a **single-stranded, negative-sense RNA virus of the Paramyxoviridae family**, the measles morbillivirus. The name itself summarises the pathology: "panencephalitis" for the diffuse involvement of both grey and white matter, "sclerosing" for the gliotic scarring the process leaves behind, and "subacute" for a tempo that is neither the days of an acute encephalitis nor the decades of a degeneration. The interval between the childhood infection and the neurological illness is the fact examiners return to, and it is precisely the fact the sources do not agree on. Kaplan and Sadock's Comprehensive Textbook of Psychiatry places it at **up to 7 to 10 years** after measles. Kaufman's Clinical Neurology for Psychiatrists frames the same interval by the patient's age: in most children and adolescents routine measles precedes SSPE by about 6 years, whereas in adults the interval stretches to about 20 years. The epidemiology explains the divergence. About 85% of patients develop SSPE in childhood, but **15%** develop it as adults, and it is this adult tail, with its far longer latency, that pulls the single "years after measles" answer apart.

TRAP **Quoting a single "years after measles" latency for SSPE as though it were fixed.** The interval genuinely differs by source and by population. Kaplan and Sadock's Comprehensive Textbook of Psychiatry and Kaufman's Clinical Neurology for Psychiatrists give different figures, and Kaufman's own account splits by age, with a short childhood interval and a much longer adult one. What a candidate should carry is not a value but the reason there is no single value: the adult-onset minority acquires measles far earlier relative to the illness, so any one figure is right only for the population it came from.

22.2 The clinical evolution and its psychiatric prodrome

The illness moves through a characteristic sequence, and its opening is the part that reaches psychiatry first. It begins with behavioural change that may last anywhere from weeks to years, accompanied by seizures and lethargy, and only later declares itself neurologically with **myoclonus**, motor and sensory disturbance, and a progressive dementia. That early behavioural phase is deceptively bland: falling school performance, irritability, apathy or a shift in personality in a child or adolescent, with nothing yet to point the clinician at the brain. As the disease advances the myoclonus that is its hallmark becomes unmistakable, jerking the child with a periodicity that mirrors the electroencephalographic discharge, and the patient eventually sinks into **akinetic mutism** before dying. The tempo from there is usually relentless and the outcome almost always fatal, and the figures are worth fixing in the mind.

12 y	1-2 y	95%
MEAN AGE AT ONSET	SURVIVAL FROM ONSET	DIE WITHIN FIVE YEARS

The disease compresses an entire dementing illness into the span of a childhood, and that speed, together with the innocuous opening, is what makes it so easy to mishandle at the start.

MNEMONIC

Behaviour, then Myoclonus, then Mutism

The order in which SSPE unfolds: a **behavioural** prodrome of weeks to years with seizures and lethargy, then the **myoclonic** phase with motor and sensory decline and dementia, ending in akinetic **mutism**. Holding the order helps, because the first stage is the one mistaken for something else.

PITFALL

Reading the early behavioural phase of SSPE as a primary psychiatric or conduct disorder. Because the behavioural change can precede the myoclonus by weeks to years, a declining, irritable adolescent may be managed as depression, oppositional behaviour or a first-episode psychosis until the first myoclonic jerks force a rethink. In a child or adolescent whose behaviour and cognition are deteriorating together, and especially where the measles or vaccination history is uncertain, the possibility should be raised early rather than after the neurology becomes obvious.

22.3 Confirming the diagnosis: cerebrospinal fluid, EEG and neuropathology

Diagnosis rests on demonstrating the immune footprint of persistent measles inside the nervous system, together with a distinctive electrical signature. The confirmatory finding is an elevated

measles antibody titre in the cerebrospinal fluid, which reflects intrathecal synthesis of antibody against the persisting virus. The point of measuring antibody in the fluid rather than in the blood is that intrathecal production signals an active infection behind the blood-brain barrier rather than mere past exposure, which is why the decisive measurement is the cerebrospinal-fluid titre read against the serum, showing antibody made within the nervous system. By contrast, an acute measles-associated neurological illness, rather than SSPE itself, can show a cerebrospinal-fluid picture of lymphocytic pleocytosis, elevated protein and positive measles IgM; in SSPE it is the markedly raised measles antibody titre, produced intrathecally, that carries the diagnosis rather than any attempt to grow the virus.

The electroencephalogram is the bedside investigation candidates are expected to recognise. It is diagnostically helpful in SSPE, classically showing **periodic sharp-wave complexes** or a burst-suppression pattern, with the periodic discharges characteristically time-locked to the myoclonic jerks. The pattern is not entirely specific, but in the right clinical setting it is strongly suggestive and it is the finding most likely to be tested.

Neuropathology completes the picture and supplies the other classic pearl. Histological examination of the brain reveals **intranuclear eosinophilic inclusions, the Cowdry bodies**, within neurons and glia, the morphological trace of the viral persistence that drives the whole illness. Three findings therefore anchor the diagnosis:

- a raised cerebrospinal-fluid measles antibody titre;
- an EEG showing periodic sharp-wave complexes or burst-suppression; and
- the Cowdry bodies on histology.

22.4 The tuberculoma: an intracerebral mass and a cause of seizures

Central-nervous-system tuberculosis most commonly presents as tuberculous meningitis, but it reaches psychiatry chiefly through the **tuberculoma**, an intracerebral mass formed by tuberculous infection. Tuberculomas are seen particularly in patients with AIDS and in people of the Indian subcontinent, and an important practical point is that computed tomography and magnetic resonance imaging can detect them even when the clinical evaluation does not suggest the diagnosis. That radiological visibility matters, because a space-occupying granuloma can sit silently until it irritates the cortex around it.

The commonest way it does so is by provoking seizures. In the Indian subcontinent tuberculomas are one of the most common causes of focal seizures, and such infections may underlie new-onset epilepsy in recent immigrants from endemic regions, so a first focal seizure in a young adult from an endemic area is a radiological question until proven otherwise. For the psychiatrist the relevance is twofold: seizures and their postictal states enter the differential of episodic behavioural disturbance, and the structural lesion itself, depending on where it lies, can contribute to focal cognitive or affective change. Where a tuberculoma sits close to eloquent or limbic cortex, the same lesion that generates the seizures can generate interictal and postictal affective and psychotic phenomena, so the structural diagnosis and the behavioural one are not separate questions but two readings of a single finding.

IN INDIA

Where tuberculosis is endemic, a new focal seizure or an unexplained ring-enhancing intracranial lesion in a young adult should keep the tuberculoma high in the differential, and cross-sectional imaging with computed tomography or magnetic resonance imaging should be obtained early rather than late. The imaging can reveal the lesion when the clinical picture does not point to tuberculosis, so the decision to image should not wait for the disease to look like tuberculosis first.

Tuberculosis of the nervous system is of course wider than the tuberculoma, and it also produces a meningitic illness with its own cerebrospinal-fluid and cranial-nerve signature. That meningitic form and its fluid profile are matters for the general-medicine and neurology texts and are not reconstructed from memory here; what belongs firmly in a psychiatric differential is the recognition that the tuberculoma is a treatable structural cause of seizures and focal deficits that imaging will find.

22.5 Whipple's disease: a systemic infection that reaches the brain

The third disorder in this group is neither viral nor mycobacterial but a systemic infection whose neurological expression is a treatable dementia. **Whipple's disease** is characterised by polyarthralgia and abdominal complaints and, in a minority of patients, by a highly distinctive movement of the eyes and jaw. Its importance to psychiatry is captured in a single phrase from the literature: it is a treatable cause to keep in mind in the demented patient. The systemic features usually precede the neurological ones by years, so the joint and gut symptoms are the clues that a cognitive decline might have a reversible infective cause. The cardinal features worth carrying are few:

- polyarthralgia, often long preceding the neurological illness;
- abdominal complaints; and
- in a minority, the pathognomonic ocular and masticatory movement described below.

The reason Whipple's disease earns a place in a psychiatry curriculum out of all proportion to its rarity is that it is one of the few dementing illnesses that can be arrested, and even reversed, once it is recognised. When a cognitive or behavioural decline is accompanied by unexplained joint and abdominal symptoms, the diagnosis should be actively considered rather than left to a specialist to raise, because the whole value of knowing about it lies in thinking of it. The formal nosology of the dementia syndrome it enters belongs to the Dementias and neurocognitive disorders notes; what this section owns is the infective cause and its one unmistakable sign.

22.6 Oculomasticatory myorhythmia: the pathognomonic sign

The sign that makes Whipple's disease examinable is **oculomasticatory myorhythmia**, and it repays precise description, because vague versions lose the marks. It consists of acquired pendular vergence oscillations of the eyes together with concurrent contraction of the muscles of mastication: a smooth, rhythmic convergence of the eyes cycling at about **1 Hz** and then diverging, accompanied by synchronous rhythmic elevation and depression of the jaw. Two

properties seal its identity. The movement persists during sleep and is unaltered by external stimuli, which together set it apart from the many oscillations that vary with attention or arousal.

The sign rarely travels alone. Patients may also show paralysis of vertical gaze, progressive somnolence and intellectual deterioration, a combination that points firmly at the brainstem and diencephalon. What gives the movement its examination value is its specificity: this distinct movement disorder has been recognised only in Whipple's disease, which is to say it is pathognomonic. A slow, rhythmic convergence of the eyes locked to a rhythmic chewing movement in a patient with cognitive decline is therefore not a sign to weigh among others but a diagnosis to act on, and it should trigger the work-up for Whipple's disease directly.

> Oculomasticatory myorhythmia, the slow pendular convergence of the eyes locked to rhythmic contraction of the muscles of mastication, has been recognised only in Whipple's disease: it is pathognomonic, and it flags a treatable dementia.

22.7 Progressive multifocal leukoencephalopathy and the wider differential

The last of the chronic neuroinfections to name here is **progressive multifocal leukoencephalopathy (PML)**, a demyelinating disease caused by the JC virus infecting oligodendrocytes. It is set beside SSPE and Whipple's disease because it too is a slowly evolving white-matter infection that can present with cognitive and behavioural change, but its clinical detail, its association with immunosuppression and its management belong with the CNS opportunistic infections of HIV and are taught there rather than reproduced here. Named alongside it, and likewise owned by other parts of these notes, are the syphilitic infection of the nervous system, taught in full with the account of neurosyphilis and general paresis, and the endemic parasitic and arboviral infections of the Indian subcontinent, taught with the endemic neuroinfections of India. Setting them side by side is useful only as a differential: the chronic infections that present to psychiatry share a habit of slow cognitive or behavioural decline, and they reward the clinician who keeps an infective cause on the list.

Within the group this section owns, three conditions are best held apart on a few shared axes, which the accompanying comparison sets out.

CONDITION	CAUSE	SIGNATURE CLINICAL CLUE	KEY INVESTIGATION
SSPE	Persistent measles morbillivirus	Myoclonus and dementia in a child or adolescent	Raised CSF measles antibody titre; periodic sharp-wave EEG
CNS tuberculoma	Tuberculosis	Focal seizures, often in an endemic-area patient	CT and MRI
Whipple's disease	Systemic infection	Oculomasticatory myorhythmia with arthralgia and gut symptoms	Directed systemic and neurological work-up

22.8 Excluding the treatable masqueraders: a clinical synthesis

What ties this heterogeneous group together is not a shared organism but a shared clinical trap and a shared reward. Each can present to psychiatry with cognitive decline, behavioural change or psychosis, and in each the diagnosis is made by the clinician who thinks of it. The stakes differ across the group. SSPE, once established, is progressive and fatal, so recognition changes prognosis and counselling rather than cure; the tuberculoma and Whipple's disease are treatable, so recognition changes the outcome outright. That asymmetry is the argument for keeping the whole set in mind, because the cost of missing the treatable members is a reversible illness allowed to run.

RULE **Before an atypical dementia or a late-onset psychosis is called primary, the treatable chronic neuroinfections must be actively excluded.** Whipple's disease is a treatable cause to keep in mind in the demented patient, and it, along with the tuberculoma and the infections owned elsewhere in these notes, will not announce itself unless it is sought. A degenerative or primary-psychiatric label is safe only once these have been thought of and set aside.

GOOD TO REMEMBER

When cognitive or behavioural decline does not fit a primary psychiatric illness or a common neurodegenerative pattern, screen deliberately for the treatable chronic infections. Look for the systemic clues of Whipple's disease, the polyarthralgia and abdominal complaints and above all the oculomasticatory myorhythmia; image early where the tuberculoma is endemic; and remember that a childhood-onset myoclonic dementia points at SSPE. The unifying action is simple: put an infective cause on the list before settling on a degenerative one.

The clinical habit these conditions teach is worth more than any single fact about them. A slowly failing intellect, a new movement disorder and a systemic clue in the history are, taken together, an invitation to look for a chronic infection of the nervous system, and the reward for looking is occasionally a patient who gets better.

Multiple sclerosis and psychoneuroimmunology

23 · Multiple sclerosis: psychiatric manifestations and the effect of disease-modifying therapy

For the psychiatrist, multiple sclerosis is one of the most psychiatrically loaded of the neurological diseases. A demyelinating illness that strikes the young adult brain, it produces mood disorder, suicide, cognitive decline and a distinctive treatment-induced psychopathology out of all proportion to its physical disability, and it does so through lesions that fall preferentially on the very white matter a stable mind depends on. The marks cluster here for reasons worth stating plainly. Depression in this disease is commoner than in comparably disabling illnesses elsewhere in medicine, its suicide risk is among the highest in clinical neurology, and its own therapies, the interferons and above all the corticosteroids, generate psychiatric syndromes that a consultant is repeatedly asked to manage. What follows sets out the

disease a psychiatrist must recognise, then works disorder by disorder through its depressive, suicidal, manic, psychotic and cognitive consequences, and closes on the psychiatric effect of disease-modifying and adjunctive treatment.

23.1 The disease a psychiatrist must recognise

Before its psychiatry can be read correctly, the disease itself has to be placed. Multiple sclerosis is a disorder of young adults, and the two standard references quote materially different central-tendency figures for its onset, a divergence worth carrying rather than smoothing over. Kaufman's Clinical Neurology for Psychiatrists gives a mean age of onset of 33 years, with 70 percent of cases arising between 21 and 40 years, and a female excess of three to four times; the Companion to Psychiatric Studies instead gives a median onset of 24 years and a female excess of about twice. Either way this is a woman's disease of early adult life, arriving at the age when identity, career and family are being built, which is part of why its psychiatric toll is so heavy.

Susceptibility is familial without being Mendelian. Compared with the general population, the disease appears 20 to 40 times more often in first-degree relatives of an affected person, with concordance of 25 percent in monozygotic and 5 percent in dizygotic twins; the Companion frames the same signal as a 30 to 50-fold familial increase. No single mutation is necessary or sufficient, so this is a polygenic, part-environmental illness rather than a single-gene disorder. Its course sorts into recognisable patterns. The relapsing-remitting form is the commonest, and here too the sources part company: Kaplan and Sadock's Comprehensive Textbook of Psychiatry puts it at 85 to 90 percent of cases while Kaufman reports it at about 80 percent at onset, with most eventually evolving to a secondary progressive phase and a smaller primary progressive minority.

FEATURE	FIGURE	SOURCE
Mean age of onset	33 years, 70% aged 21 to 40	Kaufman's Clinical Neurology for Psychiatrists
Median age of onset	24 years	Companion to Psychiatric Studies
Female predominance	3 to 4 times	Kaufman's Clinical Neurology for Psychiatrists
Female predominance	about twice	Companion to Psychiatric Studies
Relapsing-remitting share	about 80% at onset	Kaufman's Clinical Neurology for Psychiatrists
Relapsing-remitting share	85 to 90%	Kaplan and Sadock's Comprehensive Textbook

23.2 Confirming the diagnosis, and where the old teaching has moved

The diagnosis rests on demonstrating that lesions are separated in space and in time, which is the enduring core of the **McDonald criteria**. Their two chief demands have remained constant: dissemination of focal deficits in time, meaning distinct attacks separated by a period of at least a month and each lasting a minimum of 24 hours, and dissemination in space, meaning distinct

central nervous system lesions, together with the exclusion of other causes. The 2017 revision, grounded here on Kaplan and Sadock's textbook and since updated by a 2024 revision these sources do not yet carry, loosened how each may be shown: dissemination in time can now be met by a further clinical attack, by new lesions on imaging, or by the presence of central-nervous-system-specific oligoclonal bands, and dissemination in space now counts both symptomatic and asymptomatic lesions and treats cortical lesions as equivalent to juxtacortical ones.

This is exactly where a stale figure trips candidates. An older revision set a fixed imaging threshold, requiring one gadolinium-enhancing lesion or nine T2-hyperintense lesions (Kaufman); the 2017 criteria abandon the fixed nine-lesion threshold, so the frequently memorised nine-lesion rule is now obsolete and should be flagged as pre-2017 wherever it appears. Laboratory support comes from the cerebrospinal fluid, which in Kaufman's account contains **oligoclonal bands** in about 90 percent of patients or shows increased intrathecal immunoglobulin G synthesis, though these bands are nonspecific and also appear in lupus, neurosyphilis, sarcoidosis and other inflammatory disease; slowed or absent visual and somatosensory evoked potentials add support by exposing subclinical, disseminated lesions. On imaging the plaques are disseminated through the white matter but characteristically periventricular, mostly in the frontal periventricular white matter, and on sagittal magnetic resonance imaging they produce the relatively specific **Dawson fingers** radiating from the ventricles. Overall lesion load, and periventricular demyelination in particular, correlate most closely with cognitive impairment.

23.3 Depression and the psychiatric burden of the disease

Depression is the signature psychiatric consequence of multiple sclerosis and the one that carries most of the mortality. It is the most common psychiatric comorbidity of the disease, classified in DSM-5, when judged a direct physiological consequence of the disease, as a **Depressive Disorder due to another medical condition** and mapped in ICD-11 as a secondary mood syndrome; the primary nosology and pharmacotherapy of depressive illness itself belong to the Mood disorders: pharmacotherapy notes. Its frequency is striking. Major depression affects **25 to 50 percent** of patients across the lifetime, and, tellingly, it develops more often than in most other chronic, equally disabling non-neurological illnesses such as rheumatoid arthritis, which argues that the depression is not merely a reaction to being ill.

25-50%	Up to 7x	Up to 80%	45-65%
LIFETIME MAJOR DEPRESSION	SUICIDE RATE VERSUS AGE PEERS	AFFECTED BY FATIGUE OR LASSITUDE	WITH COGNITIVE DEFICITS ON TESTING

The shape of this depression is instructive. Its genetic contribution is negligible, the rate in relatives of depressed patients being considerably lower than in relatives of depressed people without the disease, and unlike primary depression it affects men and women equally. It is a disease of the brain rather than the cord: depressive symptoms are commoner when the cerebrum rather than only the spinal cord is involved, and when imaging shows cerebral atrophy and a large total lesion burden. A history of depression and premorbid trait anxiety further raise the risk.

RULE **Depression is the commonest psychiatric consequence of multiple sclerosis and exceeds the rate seen in comparably disabling illnesses elsewhere in medicine.** Because it develops more often than in chronic non-neurological disease of similar severity, and because its risk tracks cerebral rather than spinal involvement, it reflects the brain disease itself and not merely an understandable reaction to disability, and should be screened for actively.

23.4 The elevated risk of suicide

The clinical urgency of the depression is that it kills. The suicide rate among patients attending multiple sclerosis clinics runs **up to seven times** that of comparably aged people, a magnitude that largely reflects the high incidence of depression in the disease and that places this illness among the most lethal in outpatient neurology. That single figure should reset the threshold for asking about suicidal thoughts in any patient with the disease.

The risk is not spread evenly, and the profile of the patient who attempts or completes suicide is worth knowing precisely because it is partly counter-intuitive. Those affected have tended to be younger than 30 and symptomatic for less than a year, so it is the newly diagnosed rather than the long-disabled who are most exposed. Their histories carry depression in themselves or their families, alcohol misuse, and limited psychosocial support. The point examiners press hardest is the negative one: cognitive impairment from the disease is not a risk factor for suicide, so a patient's forgetfulness should never falsely reassure the clinician who is weighing risk.

MNEMONIC

Young, newly diagnosed, drinking, depressed, alone

The highest-risk profile: **young**er than 30, **newly diagnosed** and symptomatic under a year, a personal or family history of **depression**, alcohol misuse, and thin psychosocial support. Note what is absent from the list. Disease-related cognitive impairment is explicitly not a suicide risk factor, so it must not be used to downgrade concern.

23.5 Mania, euphoria and pathological affect

The elevated pole is less prominent but diagnostically richer. Bipolar disorder occurs at about twice the general-population rate in these patients, yet frank mania rarely develops, and when it does the first move is to consider a side effect of steroid treatment, which is a ready cause of a manic picture in this population. Distinct from bipolarity is **MS-induced euphoria**, an elevation of mood clearly inappropriate to the patient's disability. It is not happiness but a sign of brain disease: it tracks physical deterioration, chronicity and at least subtle intellectual impairment, is linked to steroid treatment, and usually reflects extensive cerebral involvement, so that even where psychological factors seem to explain it a diseased brain generally underlies it. It has been reported in as many as a quarter of patients, although that figure is thought to represent substantial over-reporting except in those receiving intravenous steroids, and it may relate to reduction in global grey matter volume. The classic mood-disorder eponym sometimes attached to this state is not attested in the clean reference sources and is deliberately not used here.

Closely related, and easily conflated with euphoria, is **pseudobulbar affect**, the pathological laughing and crying that is common in the disease and that in Kaufman's account is explained by pseudobulbar palsy; it may respond to serotonergic treatment. Its full phenomenology, pathophysiology and pharmacological management belong to the Stroke, TBI, delirium and focal brain syndromes notes and are named here rather than reproduced, the point for this chapter being simply that it occurs in multiple sclerosis and must be distinguished from a genuine mood disorder before either is treated.

23.6 Psychosis, and a live controversy

Psychosis is the striking exception to the disease's psychiatric reputation. On the classical teaching, and unlike depression, the prevalence of psychosis is not significantly greater in these patients than in unaffected people, and is in fact lower than in most other neurological illnesses, including Alzheimer disease, Parkinson disease, head trauma and epilepsy. Except in scattered case reports the disease does not open with psychosis, so a first psychotic episode in a young adult is far more likely to be a primary psychotic illness, whose nosology belongs to the Schizophrenia: management, antipsychotics and adverse effects notes, than a manifestation of demyelination.

TRAP **Do not state flatly that multiple sclerosis carries no excess of psychosis.** The long-standing orthodoxy, that its psychosis rate is not raised above the background population and is lower than in other neurological disease, has been challenged by a well-conducted population study reporting a modest increase concentrated in the youngest adults. The two positions genuinely disagree, and a stem can be built on either, so the examinable point is the discordance itself rather than a single number. The specific percentages from the population study are not carried in the reference sources used here and are therefore not quoted.

23.7 Cognitive impairment and fatigue

Cognitive change is common, under-recognised and characteristic in pattern. The disease produces a **subcortical dementia**, in which cortical symptoms from grey-matter injury come late if at all; memory and new learning are impaired early, while language is comparatively spared. The hallmark is slowed processing speed with impaired attention, executive deficits are common, and procedural and implicit memory are generally preserved. The most sensitive bedside-adjacent measure is the **Paced Auditory Serial Addition Test**, and characteristically the speed of test completion is more impaired than accuracy. On demanding neuropsychological testing, deficits that may be clinically silent are found in 45 to 65 percent of patients in Kaufman's account, with prevalence studies converging on 40 to 65 percent. The standardised cognitive instruments used to screen for such change are owned by the Psychotherapies, clinical assessment, MSE and nosology notes and are named here only as tools; the disease-specific point is that the deficit they detect tracks the disease's white-matter lesion burden and cerebral atrophy, in keeping with its subcortical pattern.

Overlapping with all of this is fatigue. An inexplicable, constant tiredness that neurologists call **lassitude** affects **up to 80 percent** of patients. It is entirely subjective, does not track age or degree of weakness, and intensifies both the depression and the cognitive impairment, which is

why it matters so much to the psychiatrist. Crucially, antidepressants do not relieve it, so it cannot be treated as though it were a somatic symptom of depression.

PITFALL

Reaching for an antidepressant to lift the fatigue. Because lassitude and low mood travel together, the tiredness is easily read as a depressive symptom, yet antidepressants do not alleviate the fatigue of the disease, and the reflexive prescription adds risk without benefit: selective serotonin reuptake inhibitors may increase spasticity, and antidepressants with anticholinergic action may precipitate urinary retention. Treat the depression on its own indication and address the fatigue by other means.

23.8 Treating mood, and the interferon question

Managing mood in this disease is as much about avoiding harm as about achieving benefit. Well-controlled studies have failed to prove that antidepressants improve mood in these patients, and the same agents carry the spasticity and urinary-retention hazards already noted, so an agent is chosen with these neurological adverse effects in mind. When a depression is severe enough to demand somatic treatment, the reassuring point is that the disease itself does not close off the most effective option, as the accompanying note sets out.

The disease-modifying drugs raise a separate and heavily examined question. One disease-modifying agent used in the disease is **interferon-beta**, and although early reports suggested it caused or worsened depression, that link was ultimately unsubstantiated, with recent systematic reviews finding no association; many neurologists nonetheless still avoid it in patients with a depressive history. Its near-namesake behaves entirely differently. Interferon-alpha, an innate-immune cytokine once widely used for hepatitis C and still used in some cancers, is a potent and reliable inducer of major depression in **up to 50 percent** of those treated, together with anhedonia, fatigue and cognitive impairment, and antidepressant pretreatment can lower that risk.

AGENT	TYPICAL INDICATION	DEPRESSION LINK
Interferon-beta	disease-modifying therapy for multiple sclerosis	initially reported but unsubstantiated on systematic review; cautious avoidance persists in patients with a depression history
Interferon-alpha	hepatitis C and cancer, not multiple sclerosis	a reliable inducer of major depression in up to 50 percent

TRAP **The interferon that causes depression is not the interferon used in multiple sclerosis.** The disease-modifying agent is interferon-beta, whose reported depression link is unsubstantiated on systematic review, while the reliable depressant is interferon-alpha, given for hepatitis C and cancer. A stem that offers an interferon and asks for the drug that induces major depression is testing whether the two forms have been kept apart; answer beta and you have walked into it.

GOOD TO REMEMBER

When a patient with the disease has a severe depression that needs somatic treatment, the cerebral lesions are not a contraindication to electroconvulsive therapy, which may be given with only the usual precautions. Do not let the diagnosis of multiple sclerosis, by itself, veto a treatment the depression genuinely warrants.

23.9 Corticosteroids: iatrogenic psychiatry and its management

The steroids used to treat relapses are the most reliable psychiatric offenders of all. Steroid treatment is a dependable source of psychiatric disturbance across a spectrum of syndromes, and because brain damage from the disease increases vulnerability, high-dose treatment can readily produce mental changes fulfilling DSM-5 criteria for delirium, medication-induced psychotic disorder or medication-induced depressive disorder. Glucocorticoids such as prednisone and dexamethasone are more apt to induce psychosis than mineralocorticoids such as fludrocortisone. The range runs from the mild to the severe:

- anxiety, as the mildest and commonest disturbance;
- euphoria, especially with intravenous dosing;
- mania, the picture that should prompt suspicion of a steroid effect;
- depressive symptoms, which may deepen an existing low mood;
- frank psychosis, at the severe end of the spectrum.

The dose relationship is quantified and worth carrying. The incidence of steroid psychosis, which usually begins 1 to 4 days after starting treatment, is about 4 percent in patients receiving less than 40 mg of prednisone daily and rises to 20 percent in those receiving more than 80 mg daily; once the steroid is stopped the disturbance recedes within 6 weeks in 90 percent of patients. Management follows the mechanism. If treatment cannot be interrupted, prophylactic lithium may prevent the psychosis; if it emerges, first- or second-generation antipsychotics may suppress it, whereas antidepressants may exacerbate it and should be avoided. In one reported case, **glatiramer acetate** was chosen as the disease-modifying agent considered least likely to disturb mood. Above all, because steroids are not life-saving in this disease and chronic use may raise the risk of suicidal behaviour, they should be reduced or discontinued as soon as the relapse allows.

In multiple sclerosis, depression is the commonest psychiatric burden and drives a suicide rate up to seven times that of age peers, while the disease's own treatments split cleanly: interferon-beta's depression link is unsubstantiated, whereas interferon-alpha and the corticosteroids genuinely cause psychiatric illness, so steroids are tapered as soon as the relapse allows.

24 · Psychoneuroimmunology and the immune hypothesis of psychiatric disorders

Why the marks are here. **Psychoneuroimmunology** is the study of how the immune system and the nervous system read the same threats, talk to each other, and between them shape mood,

cognition and behaviour. Two examinable ideas sit inside it. The first is the machinery of immune-brain communication: how a peripheral immune signal reaches the brain at all, and how the brain writes back to the immune system. The second is the immune or inflammatory hypothesis of psychiatric disorders, which asks whether low-grade inflammation contributes to depression and, more cautiously, to psychosis. Hold onto one framing throughout: inflammation here is a dimensional lens laid across common disorders, not a diagnostic test and not a disease in its own right. It is also distinct from the discrete antibody-mediated encephalitides, in which a single autoantibody produces a defined neuropsychiatric syndrome; psychoneuroimmunology is instead about graded, subthreshold inflammation nudging the risk and the colour of otherwise common disorders. That framing is also where candidates lose marks, because the biomarker data are real but far less specific than they first appear.

RULE **Immune-brain signalling runs both ways.** The immune system can shape behaviour to promote healing and recovery, and neuronal activity can in turn alter the immune response inside the brain. Any model that treats the brain as merely downstream of a peripheral immune signal, or the immune system as a passive effector of central commands, has drawn only half the circuit.

24.1 Two systems, one conversation

The immune and nervous systems are the body's two great surveillance-and-response networks, and psychoneuroimmunology begins by noticing how alike they are. Both detect and neutralise threats and work to maintain homeostasis; both perceive the internal and external environment and deploy specific receptors to monitor ambient stimuli. What the field adds is that the two are coupled rather than parallel. On one side, the immune system can control behaviour in a manner that promotes healing and recovery, the phenomenon that later becomes sickness behaviour. On the other, neuronal activity can alter the immunologic response within the brain, with microglia sensing spillover neurotransmitters through their surface receptors and adjusting their state accordingly. So the traffic is genuinely bidirectional, and the rest of this topic simply follows the signal in each direction: outward from a stressed brain to the marrow and lymphoid organs, and inward from an inflamed body to mood and thought. The symmetry is not merely rhetorical; it means a brain state can be read off the immune compartment, and an immune state can be read off behaviour, which is exactly what the biomarker and sickness-behaviour literatures later exploit. Keeping the loop in view matters clinically, because it suggests that addressing only one arm, calming the mind while inflammation smoulders, or damping cytokines while chronic stress persists, may fall short, an idea that motivates combined approaches rather than a proven treatment rule.

24.2 How the immune signal reaches the brain

The brain was long taught to be immune-privileged; the modern account is regulated access, not exclusion. Resident immune cells, the microglia, sit mostly in the parenchyma, while trafficking immune cells are stationed in specialised niches, chiefly the meninges and the choroid plexus. The blood-brain barrier tightly regulates the flow of peripheral immune molecules into the central nervous system, and the usual site of blood-brain interchange is the fenestrated capillaries

of the **circumventricular organs**. A further conduit is the **glymphatic system**, the brain's lymphatic route, which carries cytokines from the meningeal compartment through cerebrospinal fluid and the perivascular (Virchow-Robin) spaces into the parenchyma. It is worth reducing all of this to three routes, which the accompanying figure sets out: a humoral route for soluble mediators, a cellular route for trafficking leukocytes, and a neural route that the next subtopic develops.

ROUTE	WHERE IT OPERATES	WHAT IS CARRIED
Humoral	Fenestrated capillaries of the circumventricular organs; glymphatic and perivascular flow	Soluble cytokines from the periphery and meningeal compartment
Cellular	Meninges and choroid plexus niches	Trafficking immune cells sampling the interface
Neural	Vagal afferents and dorsal root ganglion sensory neurons	Immunologic signals read by Toll-like and cytokine receptors

24.3 The neural and autonomic interface

The neural route deserves its own account because it is fast and specific. Both the sensory afferent fibres of the autonomic nervous system, the vagal afferents, and the peripheral sensory neurons of the dorsal root ganglia display immune receptors, namely Toll-like receptors and cytokine receptors, and they transmit immunologic signals to the brain. When these receptors are stimulated the neurons become more sensitive to pain and itch, which drives behavioural change such as avoidance, so the afferent arm is not just a wire but a behaviour-shaping input in its own right. Traffic in the other direction is autonomic. Sympathetic nervous system fibres innervate the reticuloendothelial organs, the bone marrow, thymus, spleen and lymph nodes, while efferent input from the vagal and recurrent laryngeal nerves supplies parasympathetic counter-regulation. The immune cells themselves are wired to listen: the adrenergic receptor on T and B lymphocytes was one of the first identified, followed by receptors for noradrenaline, dopamine, acetylcholine, serotonin and glutamate on immune cells including microglia. Timing is everything with these mediators. Norepinephrine boosts immune activity during antigen recognition but suppresses it during the effector phase, so the same signal helps or hinders depending on when, and at what dose, the cell meets it. The practical corollary is that the autonomic state a patient inhabits, sympathetically aroused or parasympathetically settled, biases the immune response they can mount, which is one concrete channel through which sustained stress leaves an immune signature.

24.4 From stress to inflammation

The cascade that turns a stressor into measurable inflammation is the mechanistic hinge of the whole topic. Noradrenaline released by activated sympathetic fibres stimulates the bone marrow to produce myeloid cells. Endogenous danger signals and microbial signals, the DAMPs and MAMPs, then switch on two inflammatory pathways: NF-kappa-B and the **NLRP3 inflammasome**. The inflammasome activates caspase-1, which matures IL-1-beta and IL-18 and cleaves the glucocorticoid receptor, contributing to **glucocorticoid resistance**, the state in which

immune cells stop heeding cortisol's usual restraining signal. NF-kappa-B, for its part, drives the release of TNF and IL-6. These cytokines, IL-1-beta, IL-18, TNF and IL-6, then access the brain by the humoral and neural routes already described. This is why a psychological stressor can end as a rise in circulating IL-6 and TNF, and why chronic stress and depression are so often associated with a blunted cortisol brake rather than a simple excess or deficiency of cortisol. The link back to the hypothalamic-pituitary-adrenal axis is direct: a receptor that has been cleaved cannot transmit feedback.

24.5 Acute versus chronic stress

Stress does not do one thing to immunity; its direction depends on duration, and confusing the two is a common conceptual error. Acute stress boosts leukocyte activity, recruiting cells from the bone marrow and driving infiltration and proliferation. That acute enhancement is a sympathetic effect, and the suppression that follows once the stressor resolves is a parasympathetic one. Chronic or severe stress does the reverse and more. It shifts the balance away from proinflammatory Th1 towards anti-inflammatory Th2, suppressing interferon-gamma and IL-2; it lowers natural killer cell activity and lymphocyte proliferation; and it decreases the sensitivity of monocytes to endogenous cortisol, which loops straight back into glucocorticoid resistance. The clinical reading is that acute and chronic stress are not two points on one line but qualitatively different immune states, which is why a brief challenge and a grinding adversity produce different immune and behavioural pictures. It also resolves an apparent paradox, that stress is described as both proinflammatory and immunosuppressive at once. The acute, sympathetically driven phase amplifies innate inflammatory activity, while the sustained phase suppresses cell-mediated defence through the Th1 to Th2 shift and falling interferon-gamma, yet simultaneously leaves the innate inflammatory tone poorly restrained because monocytes no longer respond fully to cortisol. The two halves are sequential rather than contradictory, and depression tends to sit at the chronic end, where the cortisol brake has failed.

DIMENSION	ACUTE STRESS	CHRONIC OR SEVERE STRESS
Net effect on immunity	Enhancement	Suppression and dysregulation
Leukocytes	Recruitment, infiltration, proliferation	Reduced lymphocyte proliferation
Natural killer cells	Not the primary change	Decreased activity
T-helper balance	Not the primary change	Th1 to Th2 shift; IFN-gamma and IL-2 suppressed
Autonomic basis	Sympathetic drive, then parasympathetic suppression on resolution	Sustained sympathetic and HPA drive with glucocorticoid resistance
Cortisol sensitivity	Not the primary change	Monocytes less sensitive

24.6 Sickness behaviour: the bridge to depression

Everything so far becomes clinically legible through **sickness behaviours**. Coincident with the physiologic immune changes, animals given an immune challenge reduce a recognisable cluster of

behaviours.

- reduced general activity;
- reduced exploration of novel objects;
- reduced social interaction;
- reduced food and water intake; and
- reduced sexual behaviour.

The same link appears in humans. Healthy volunteers develop anxiety and depression in response to low-dose cytokines, and even people with the common cold show sickness behaviours such as fatigue, loss of interest, low mood and cognitive impairment. Tricyclic antidepressants reduce lipopolysaccharide-induced sickness behaviour in animals, an early pharmacological hint that this pathway overlaps with the pathway of depression. The concept earns its place because it is the translational bridge: it explains how a proinflammatory signal is converted into a depressive phenotype, and it reframes the vegetative symptoms of depression as an evolved, cytokine-driven programme rather than an incidental accompaniment of low mood.

PITFALL

The overlap cuts both ways at the bedside. The fatigue, anhedonia, poor concentration and low mood of an acute systemic infection can closely mimic a depressive episode, because sickness behaviour and the vegetative symptoms of depression share a cytokine-driven substrate. Reading transient sickness behaviour as a new primary depression, and starting an antidepressant on the strength of it, is a real and avoidable error; a picture that lifts as the infection clears usually reflects transient sickness behaviour rather than a primary mood disorder, though the judgement rests on the full syndrome and course.

24.7 How cytokines produce depression

If inflammation contributes to depression, it does so through three converging mechanisms, and holding all three together is what the examiners reward. Proinflammatory cytokines, including the type I and type II interferons, IL-1-beta and TNF, act on the monoaminergic, tryptophan and glutamatergic systems at once. First, they reduce the availability of the monoamines serotonin, dopamine and noradrenaline, both by increasing the expression and function of presynaptic reuptake transporters through MAPK pathways and by depleting tetrahydrobiopterin, the cofactor for monoamine synthesis. Second, they activate **indoleamine 2,3-dioxygenase**, the enzyme that diverts tryptophan down the kynurenine pathway; microglia then convert kynurenine to **quinolinic acid**, which activates extrasynaptic NMDA receptors. Third, together with a cytokine-induced fall in astrocytic glutamate reuptake, the result is excess glutamate, reduced brain-derived neurotrophic factor and excitotoxicity. The elegance for revision is that a single upstream signal, inflammation, converges on the monoaminergic, tryptophan and glutamatergic systems at once, overlapping the monoaminergic and glutamatergic targets that existing antidepressants act on, while the kynurenine arm remains mainly an investigational one.

MNEMONIC

Starve the monoamines, hijack the tryptophan, flood the synapse

Three routes from cytokine to low mood: monoamines are **starved** by raised reuptake and lost tetrahydrobiopterin; tryptophan is **hijacked** down the kynurenine pathway to quinolinic acid; and the synapse is **flooded** with glutamate as astrocytic reuptake falls, lowering BDNF and driving excitotoxicity.

24.8 Inflammatory biomarkers and longitudinal risk in depression

From mechanism to measurement. Individuals with depressive disorders show elevated levels of several pro- and anti-inflammatory cytokines, including IL-2, IL-5, IL-6, IL-7, IL-8, IL-10, IL-12, IL-13, interferon-gamma, C-reactive protein and TNF. From that broad set, the most robust and consistent associations are with CRP, IL-6, IL-1 and TNF. A meta-analysis of 82 studies found peripheral IL-6 and TNF-alpha elevated and interferon-gamma lower in major depression, with TNF-alpha, IL-6 and IL-1b tracking the severity of depressive and cognitive symptoms. The relationship is prospective as well as cross-sectional: raised CRP and IL-6 predict later depression, and in a large Danish population cohort the risk of mood disorder rose by **45%** after a diagnosis of autoimmune disorder and by **62%** after hospitalisation for an infectious disease. IL-6 carries particular weight: it predicts the chronicity and follow-up severity of depression and is associated with increased suicide risk, the highest levels correlating with the most violent suicide attempts, and autopsies of suicide victims show raised cytokine concentrations in the orbitofrontal and dorsolateral prefrontal cortex. That combination, a prospective marker of chronicity and a correlate of violent self-harm with a tissue-level footprint, is why IL-6 recurs across the whole immune literature rather than sitting as one cytokine among many. None of this makes inflammation the diagnosis. The primary nosology and management of depression belong to the Mood disorders: pharmacotherapy notes; what the immune lens adds is an inflammatory subgroup and a set of state markers.

GOOD TO REMEMBER

In depression that is proving treatment-resistant, a coexisting autoimmune or infective illness, or a prospective rise in CRP and IL-6, is worth noting as a group-level correlate of severity and chronicity and a prompt to look for a treatable medical contributor. It does not by itself establish causation or a diagnosis, but it flags the patient in whom the mood disorder and the body are moving together.

24.9 The immune lens on schizophrenia

The immune story is weaker and more cautious in psychosis than in depression, and that caution is itself examinable. Increased inflammation is generally not associated with the development of positive psychotic symptoms, and circulating immune markers do not distinguish schizophrenia from affective disorders. What the data do support is modest. A meta-analysis of over 85,000 people found small but reliable increases in CRP among people with schizophrenia, paralleling positive but not negative symptom severity, and a first-episode meta-analysis comparing 3,453

drug-naive first-episode patients with 4,152 controls reported raised IL-6, IL-8 and TNF. Elevated IL-6, TNF, soluble IL-2 receptor and IL-1 receptor antagonist appear across acute schizophrenia, bipolar mania and major depression alike, which underlines the non-specificity. Two further strands complete the picture. A Danish National Registry analysis linked prior illness to later schizophrenia, and the developmental hypothesis of **maternal immune activation** ties prenatal rubella, influenza and toxoplasmosis to schizophrenia; maternal IL-8 is significantly raised in pregnancies that lead to schizophrenia, with fetal exposure associated with increased ventricular cerebrospinal fluid and reduced left entorhinal cortex and right posterior cingulate volumes, raised maternal TNF has been linked to later illness, and blocking IL-1, IL-6 and TNF reduces schizophrenia-like deficits in animal models. As with mood disorder, the primary nosology and treatment sit in the Schizophrenia: management, antipsychotics and adverse effects notes.

29%	60%	125%
HIGHER SCHIZOPHRENIA RISK AFTER AUTOIMMUNE DISORDER	AFTER HOSPITALISATION FOR INFECTION	WHEN AUTOIMMUNITY AND INFECTION COMBINE

TRAP **Raised inflammatory markers are not a psychosis-specific finding.** Candidates read an elevated CRP or IL-6 as evidence that inflammation is driving the hallucinations and delusions, or as a marker that separates schizophrenia from mood disorder. Neither holds: inflammation is generally not tied to the development of positive psychotic symptoms, and the same circulating markers rise across acute schizophrenia, bipolar mania and depression. The lens describes a shared vulnerability, not a diagnostic signature.

IN INDIA

When a new depressive or psychotic presentation coincides with a systemic infection or an autoimmune illness, the raised prospective risk shown above is a reason to take a deliberate infective and autoimmune history and to look for an active inflammatory process before settling on a primary psychiatric diagnosis. That discipline earns its keep wherever infective and autoimmune illness routinely share the psychiatric differential, and it keeps a treatable contributor from being missed.

Immune-brain signalling is bidirectional, and subthreshold inflammation is both a state marker and, in some patients, a probable contributor across mood and psychotic illness: a lens on severity and vulnerability, never a diagnosis in itself.

Degenerative, neuromuscular and functional neurology

25 · Atypical parkinsonism and the hereditary ataxias

Two families of neurodegeneration bring a patient to psychiatric attention long before anyone has thought of a psychiatrist. The first is the group of atypical parkinsonian, or Parkinson-plus, syndromes, which borrow the outward signs of Parkinson disease but run a faster

course, resist dopaminergic treatment, and carry cognitive and behavioural change from early in the illness. The second is the hereditary cerebellar ataxias, in which a genetically determined degeneration of the cerebellum and its connections produces not only incoordination but a substantial and frequently unrecognised psychiatric burden. Marks cluster here because each of these disorders is first misread as something else: a Parkinson-plus syndrome as ordinary Parkinson disease, an ataxia as a primary depression or a personality change. Recognising the atypical features, and knowing which signs separate one syndrome from the next, is what the examination rewards.

25.1 The Parkinson-plus class and where it departs from Parkinson disease

Parkinson-plus diseases are a group of related neurodegenerative illnesses that share many of the physical signs of Parkinson disease, yet other features set them apart from it and from one another. The single most useful class-level discriminator is that, compared with Parkinson disease, they follow a more rapid course and respond less well to dopaminergic medication. A patient whose parkinsonism advances quickly, spreads symmetrically, and fails to improve convincingly on levodopa is telling you the diagnosis is probably not idiopathic Parkinson disease.

A second early signpost is the timing of dementia. In several illnesses characterised by parkinsonism, among them dementia pugilistica, the Parkinson-plus disorders themselves, and dementia with Lewy bodies, dementia may appear as the first or the most prominent symptom, whereas in Parkinson disease cognitive decline is a comparatively late development. Dementia with Lewy bodies, the one-year rule that separates it from Parkinson disease dementia, and its neuroleptic sensitivity are taught in the Dementias and neurocognitive disorders notes and are only named here. The point to carry forward is that early, prominent cognitive or behavioural change in a parkinsonian patient argues against simple Parkinson disease.

RULE **A parkinsonism that progresses rapidly, carries early atypical features, and responds poorly to levodopa is a Parkinson-plus syndrome until proven otherwise.** Idiopathic Parkinson disease is asymmetric at onset, evolves slowly, and answers to dopaminergic treatment; symmetry at onset is not a class-level rule but varies by syndrome, being typical of progressive supranuclear palsy yet absent in the characteristically asymmetric corticobasal degeneration. Each of the three degenerative Parkinson-plus syndromes then adds a signature of its own: a vertical gaze palsy, autonomic or cerebellar failure, or an apraxic alien limb.

25.2 Progressive supranuclear palsy: axial rigidity, early falls and the vertical gaze palsy

Progressive supranuclear palsy presents with parkinsonism but rarely with the resting tremor that typifies Parkinson disease. Its rigidity is predominantly axial, forcing the head, neck and spine to remain overly upright and unnaturally straight, in pointed contrast to the flexed neck and stooped kyphosis of Parkinson disease. Postural instability and falls occur early, often within the first year of onset, and the parkinsonism progresses with symmetric bradykinesia and axial rigidity rather than the asymmetric picture of Parkinson disease. It becomes manifest in affected individuals over the age of 40 years, and levodopa does not significantly correct it.

The most characteristic feature is a loss of the ability to look voluntarily in the vertical plane, the **vertical supranuclear gaze palsy**, which characteristically impairs downward gaze first and can make descending stairs treacherous. It is named supranuclear for a reason that also supplies the bedside test: the lesion sits above the ocular motor nuclei, so the oculocephalic, or doll's eye, vertical reflex is preserved even when voluntary vertical gaze is lost. This sign may not appear until **3 years** after the parkinsonism begins, so its absence early does not exclude the diagnosis. Progressive supranuclear palsy is a tauopathy, and MRI may show the hummingbird sign of rostral midbrain atrophy.

MNEMONIC

Look down, stand straight, fall back

The clinical spine of progressive supranuclear palsy: impaired voluntary **down**gaze, an unnaturally **straight** and upright axial posture, and early **backward** falls from postural instability, none of which behaves like idiopathic Parkinson disease.

PITFALL

Being reassured by intact reflex eye movements. Because the deficit is supranuclear, the doll's eye vertical reflex is preserved while *voluntary* vertical gaze is lost, so the eyes can look normal on passive testing. Ask the patient to look up and down to command; a normal oculocephalic response does not exclude the palsy.

25.3 The subcortical dementia and neuropsychiatry of progressive supranuclear palsy

Cognitive change in progressive supranuclear palsy is common and follows a frontal-subcortical pattern rather than the cortical one of Alzheimer disease. Roughly 10% of patients present with cognitive symptoms, and up to **70%** develop a frank dementia over the course of the illness. Neurocognitive assessment shows psychomotor slowing, poor working memory and executive dysfunction, the profile of a disease that disrupts the frontal circuits from below. The classic cortical triad of aphasia, apraxia and agnosia is not the story here, which is one of the features that separates it from Alzheimer disease and, as the next sections show, from corticobasal degeneration.

Progressive supranuclear palsy also produces pseudobulbar signs, the disinhibition of brainstem motor programmes that in its affective form yields pathological laughing and crying. That syndrome in full, its phenomenology, pathophysiology and treatment, is taught in the Stroke, TBI, delirium and focal brain syndromes notes and is named here only as one of the neuropsychiatric accompaniments of the disease. Retropulsion, the tendency to lose balance in a backward direction, is often prominent and compounds the falls that already mark the illness from its onset. Taken together, the slowing, the executive failure and the brainstem release give the disease a characteristic frontal-subcortical texture that becomes diagnostically useful when it is set against the cortical picture of corticobasal degeneration.

25.4 Multiple system atrophy: two subtypes and the levodopa wall

Multiple system atrophy is a family of Parkinson-plus disease with two classic subtypes a candidate is expected to name. One is a cerebellar subtype, historically called olivopontocerebellar degeneration, notable for ataxia; the other is a subtype dominated by pronounced autonomic dysfunction, historically called Shy-Drager syndrome. Naming the disease by the system that fails first, cerebellum or autonomic nervous system, is the practical way to hold its presentations together.

Like Parkinson disease, multiple system atrophy is a synucleinopathy; unlike Parkinson disease, it responds poorly to levodopa, and any response is usually incomplete or unsustained, which is a large part of what marks it as atypical. The autonomic failure of the dysautonomic subtype expresses itself as orthostatic intolerance, disturbed bladder control and other signs of a nervous system that can no longer regulate its own housekeeping, and it is frequently the autonomic collapse, rather than the parkinsonism, that disables the patient and brings them to attention. The cerebellar subtype, by contrast, stands at the border of the second half of this section, because its ataxia raises exactly the differential that the hereditary cerebellar degenerations occupy, and a middle-aged patient with a progressive ataxia plus autonomic failure and a levodopa-resistant parkinsonism is describing this disease.

25.5 Corticobasal degeneration: the asymmetric apraxic limb

Corticobasal degeneration has as its core clinical syndrome a progressive asymmetric rigidity and apraxia together with signs of both cortical and basal-ganglia dysfunction. It was first identified in 1967 in three individuals with a progressive asymmetrical akinetic-rigid syndrome and apraxia. It is largely sporadic and, like progressive supranuclear palsy, is a tauopathy; more precisely it is a sporadic four-repeat tauopathy, with deposition of hyperphosphorylated four-repeat tau in neurons and glia as astrocytic plaques. Because the living clinical picture and the histopathology do not always coincide, the term corticobasal syndrome is often used for the clinical presentation, reserving the label corticobasal degeneration for the pathological diagnosis.

The clinical signs cluster on one side of the body and betray the cortex as much as the basal ganglia:

- asymmetric rigidity and limb apraxia;
- the **alien limb phenomenon**, in which a limb behaves as if it has a will of its own;
- cortical sensory loss;
- myoclonus; and
- postural instability.

The neuropathology matches the asymmetry of the signs. Imaging shows asymmetric frontoparietal atrophy with subcortical hypometabolism contralateral to the more clinically affected side, and tau inclusions are found in the substantia nigra, the basal ganglia and along the dentatorubrothalamic tracts. The laterality is not incidental; it is the anatomical signature of a

disease that attacks one hemisphere ahead of the other, and it mirrors the one-sided apraxia and alien limb at the bedside.

25.6 The neuropsychiatry of corticobasal degeneration and its contrast with progressive supranuclear palsy

The neuropsychiatric hallmarks of corticobasal syndrome are cortical in character, which is what most sharply separates it from progressive supranuclear palsy. Its clinical signatures are a lateralised ideomotor and ideational apraxia and the alien hand syndrome, and on cognitive testing parietal-lobe dysfunction shows itself as sensory or visual neglect and visuospatial impairment, alongside measurable deficits in social cognition. Where progressive supranuclear palsy disrupts the frontal circuits from below and produces a slowed, subcortical picture, corticobasal degeneration attacks the parietal and frontal cortex directly and produces the apraxic, neglecting profile of a cortical disease. That localising difference is the deeper reason the two tauopathies feel so unlike each other at the bedside despite sharing a four-repeat tau pathology.

GOOD TO REMEMBER

When two akinetic-rigid, levodopa-poor patients differ in mood, the affective colour can support the distinction the motor signs already draw: depression is much more common in corticobasal degeneration, whereas apathy is the hallmark of progressive supranuclear palsy. It corroborates rather than leads; the side-to-side asymmetry and apraxia of the one and the axial, gaze-palsied posture of the other remain the diagnostic anchors, with the affective colour a genuine supporting clue rather than an incidental comorbidity.

Set side by side, the three degenerative Parkinson-plus syndromes separate cleanly on a handful of axes, as summarised in the accompanying comparison.

FEATURE	PSP	MSA	CBD
Signature motor sign	Axial rigidity, early falls	Ataxia or autonomic failure	Asymmetric akinesia, apraxia
Giveaway sign	Vertical supranuclear gaze palsy	Autonomic or cerebellar failure	Alien limb, cortical sensory loss
Symmetry	Symmetric	-	Asymmetric
Levodopa	Poor response	Poor or unsustained response	Poor response
Pathology	Tauopathy	Synucleinopathy	Four-repeat tauopathy

TRAP **Pushing the levodopa dose ever higher when a "Parkinson disease" patient will not respond.** In the Parkinson-plus syndromes the poor response to dopaminergic medication is not undertreatment, it is the diagnosis. Multiple system atrophy responds poorly, if at all, to levodopa, and progressive supranuclear palsy responds poorly; a parkinsonism that will not answer escalating doses should prompt a rethink of the diagnosis, not a bigger prescription.

25.7 Medication-induced parkinsonism: the mimic on the prescription pad

Medication-induced parkinsonism is the most common and most reversible of the Parkinson disease mimics, and it belongs in every differential because its cause sits in the drug chart. Typical and most atypical antipsychotic agents, to a greater or lesser degree, block D2 receptors and so routinely cause parkinsonism. So do several non-psychiatric D2 blockers prescribed for nausea and vertigo:

- metoclopramide;
- prochlorperazine; and
- promethazine.

Tetrabenazine produces the same picture by a different route, depleting dopamine rather than blocking its receptor. The resemblance to idiopathic Parkinson disease is close enough that clinical examination cannot reliably tell them apart; the most useful single clue is symmetry. Medication-induced parkinsonism usually causes symmetric, bilateral signs from the onset, whereas Parkinson disease tends to begin asymmetrically and stays asymmetric throughout. Once the offending drug is stopped the syndrome usually resolves over weeks, but a minority does not: about 10% of persistent cases turn out to harbour Parkinson disease or a Parkinson-plus syndrome that the medication had merely unmasked.

> **IN INDIA**
>
> In any patient presenting with new, symmetric, bilateral parkinsonism, review the drug chart before diagnosing Parkinson disease. D2-blocking antiemetics such as metoclopramide and prochlorperazine, and antipsychotics of every class, are common precipitants; withdraw the offending agent and reassess over the following weeks. The mimic is common and usually reversible, and the price of missing it is months of unnecessary antiparkinsonian treatment for a problem the prescription created.

25.8 The hereditary ataxias: inheritance, prevalence and age of onset

The second half of this section turns from the basal ganglia to the cerebellum. The hereditary ataxias are genetically heterogeneous, but an important group of them, including Friedreich ataxia and many spinocerebellar ataxias, are trinucleotide-repeat disorders, and the first thing to fix is the pattern of inheritance. **Friedreich ataxia** follows an autosomal recessive pattern and is the most common autosomal recessive cerebellar ataxia, while the **spinocerebellar ataxias** follow an autosomal dominant pattern; the same trinucleotide-repeat mechanism, working in different genes, also underlies the autosomal dominant Huntington disease and the sex-linked fragile X syndrome, which are named here only to place the ataxias within their family. The spinocerebellar ataxias themselves form a large numbered series, running to types 1 to 42 and beyond.

Prevalence gives a second anchor. Taken separately, the recessive cerebellar degenerations (primarily Friedreich ataxia), the dominant ones, and the sporadic cerebellar neurodegenerations each run at about **2 per 100,000**. Age of onset then separates the two hereditary groups at the

bedside: Friedreich ataxia declares itself in childhood, whereas the spinocerebellar ataxias typically begin in middle age. That single discriminator, a child with a progressive ataxia against an adult with one, does much of the early diagnostic work before any genetic test is sent, and it maps onto the recessive-versus-dominant split that structures the whole group. The specific repeat unit and causative gene of each disorder belong with the individual disease and are not asserted where the available sources do not cleanly supply them.

25.9 The psychiatric burden of the cerebellar degenerations

For the psychiatrist the important truth about these disorders is that the cerebellum is not merely a motor structure, and its degeneration carries a heavy and often silent psychiatric load. In degenerative disorders predominantly limited to the cerebellum, most patients develop one or more psychiatric syndromes within the first decade and a half of the illness, and depression dominates that picture.

about 75%	about one-third	about 25%
DEVELOP A PSYCHIATRIC SYNDROME WITHIN 10 TO 15 YEARS OF ONSET	HAVE AN EPISODE OF MAJOR DEPRESSION	DEVELOP PERSISTENT PERSONALITY CHANGE

Beyond the patients who reach the threshold for major depression, roughly as many again experience other depressive states, from brief recurrent and minor depression to dysthymia, so that depressive illness of one form or another comes close to being the rule. The persistent personality change resembles that seen in Huntington disease and, as in Huntington disease, is not confined to those who have gone on to cognitive decline. Where extra-cerebellar structures degenerate as well, the neuropsychiatric picture widens: **spinocerebellar ataxia type 2**, a late-onset autosomal dominant disorder in which loss extends to the inferior olive, pontocerebellar nuclei, substantia nigra and motor neurons, adds executive dysfunction, impaired memory, apathy, depression and potentially psychosis to the cerebellar signs.

Suicidality can be striking, and one figure is worth carrying. In **spinocerebellar ataxia type 3**, also called Machado-Joseph disease, the rate of suicidal ideation has been reported as high as **65%**. That number, from The Clinical Research Consortium for Spinocerebellar Ataxias, is specific to type 3 and stated as an upper bound rather than a general figure for the ataxias, but it is a stark reminder that these are not benign motor disorders.

TRAP **Waiting for a cerebellar-degeneration patient to complain of low mood.** Depression is common in these disorders, yet self-awareness of it is frequently absent, so a clinician who screens only by asking the patient will miss it. Depression here must be sought actively and corroborated with an informant, and the high figures for suicidal ideation make that active screening a safety matter rather than a courtesy.

Atypical parkinsonism is a rapid, symmetric, levodopa-poor parkinsonism with a signature sign for each disease, a vertical gaze palsy in progressive supranuclear palsy, autonomic or cerebellar failure in multiple system atrophy, an apraxic alien limb in corticobasal degeneration, while the hereditary ataxias carry a heavy and frequently unrecognised psychiatric burden that must be sought and not awaited.

26 · Motor neuron disease: cognitive, behavioural and end-of-life neuropsychiatry

Motor neuron disease was for a long time taught as a pure disorder of movement, and that is exactly why its psychiatry catches candidates out. The illness that supposedly spares the mind is now understood to reach the frontal and temporal lobes, to overlap with a dementia syndrome, to generate a disinhibited disorder of affective display, and to force some of the hardest end-of-life decisions in clinical medicine onto a patient whose intellect may itself be quietly eroding. The marks here reward the psychiatrist who can hold two facts at once: that the motor picture is the diagnosis, and that cognition, behaviour and mood belong to the same disease and shape everything from capacity to survival. What follows sets out the disease a psychiatrist must recognise, then works through its cognitive overlap with frontotemporal degeneration, its genetic and behavioural links, the pathological affect it produces, the drugs that modify and palliate it, and the prognostic frame within which every psychiatric decision has to be made.

26.1 The disease a psychiatrist must recognise

The commonest form of motor neuron disease is **amyotrophic lateral sclerosis**, long regarded as a very rare condition though its frequency is far from negligible. On the account in Kaplan and Sadock's Comprehensive Textbook of Psychiatry it is uncommon rather than vanishingly rare, a distinction that matters because a psychiatrist working a busy neurology liaison service will meet it. Its diagnostic signature is the coexistence, in the same patient, of two families of physical sign that are usually thought of as opposites. The lower motor neuron signs are weakness, muscle atrophy and fasciculations, the last being **fasciculations**, painless spontaneous discharges of motor units seen as rapid involuntary twitches in limb or trunk muscles, while the upper motor neuron signs are spasticity, loss of dexterity, hyperreflexia and pseudobulbar palsy. It is that combination, upper and lower motor neuron signs together, that classically makes the clinical diagnosis.

SIGN CLASS	CLINICAL FEATURES
Lower motor neuron	weakness, muscle atrophy and fasciculations (painless spontaneous discharges of motor units, seen as involuntary twitches in limb or trunk muscles)
Upper motor neuron	spasticity, loss of dexterity, hyperreflexia and pseudobulbar palsy

~2 / 100,000	6 / 100,000	5-20%	within 3 years
ANNUAL INCIDENCE	POINT PREVALENCE	FAMILIAL CASES	TYPICAL SURVIVAL

26.2 Where the degeneration falls

The psychiatry of this disease is legible only once its pathology is placed, because the lesions that produce the mind symptoms are not where the classical teaching looks. The degeneration is multi-level: it consumes the motor neurons of the spinal cord, that is the anterior horn cells, together with the motor neurons of the brainstem and the upper motor neurons of the cerebrum. So far this is the map of a movement disorder. The critical addition for a psychiatrist is cortical. In the same disease there is gliosis and neuronal loss in the superficial layers of the dorsomedial neocortex of the frontal lobes, and this may extend into the hippocampal and parahippocampal regions of the temporal lobes.

That frontal and temporal cortical involvement is not incidental. It is the anatomical substrate of everything the rest of this topic describes: the executive and language changes, the behavioural syndromes, and the overlap with a frontotemporal dementia. A disease whose pathology reaches the dorsomedial frontal neocortex and the medial temporal structures is, on anatomical grounds alone, a disease that should be expected to disturb cognition and social behaviour, and the reason the older teaching missed this is simply that the motor signs are so much louder that the cortical ones went unexamined. Reading the pathology this way turns the neuropsychiatric findings from a surprise into a prediction.

26.3 The amyotrophic lateral sclerosis and frontotemporal degeneration spectrum

The single most important conceptual shift in this field is that **frontotemporal degeneration** and amyotrophic lateral sclerosis lie on one spectrum rather than standing as two wholly separate diseases. There is increasing recognition of the neuropsychiatric aspects of the motor disease, most notably its association with frontotemporal degeneration, and careful examination shows cognitive decline in patients who present with otherwise classic motor symptoms; either condition may develop first, so the dementia may emerge in someone with established motor disease or the motor disease in someone already carrying the dementia diagnosis. The two share much of their molecular pathology, the ubiquitin-positive neuronal cytoplasmic inclusions found in lower motor neurons, hippocampus and neocortex, staining in most cases for the protein TDP-43, the pathological thread that lets them be understood as one disease spectrum rather than two unrelated illnesses. The standalone account of frontotemporal dementia as a primary neurocognitive disorder belongs to the Dementias and neurocognitive disorders notes and is named here, not reproduced.

Quantified, the overlap is substantial rather than anecdotal. Between **one-third to one-half** of patients perform poorly on formal neuropsychological testing against matched controls, and the pattern is a frontal one, distinct from the amnestic and apraxic profile of Alzheimer disease. Its recognisable components are worth carrying:

- mild memory loss, rather than the dense amnesia of a cortical dementia;
- word-finding difficulty and reduced semantic fluency;
- executive dysfunction, the frontal signature of the disease;

- impaired recognition of facial emotion, a social-cognition deficit that patients and families feel long before it is tested.

The reason this matters clinically, and not only diagnostically, is prognostic: the presence of cognitive impairment is associated with shorter survival. The standardised cognitive instruments used to detect such change are owned by the Psychotherapies, clinical assessment, MSE and nosology notes and are named here only as tools.

RULE **Motor neuron disease is not a pure disorder of movement.** It lies on a continuum with frontotemporal degeneration, and a substantial minority of patients carry a frontal-executive cognitive syndrome that the loud motor signs conceal. Because that impairment predicts a shorter survival and erodes the very faculties a patient needs to make end-of-life decisions, cognition and behaviour must be assessed as part of the diagnosis, not left until they declare themselves.

GOOD TO REMEMBER

Screen cognition and behaviour at diagnosis, before the motor decline dominates the consultation. Doing so is not academic: because cognitive impairment is linked to decreased survival time and may impair a patient's capacity to weigh decisions about ventilation and feeding, though such impairment does not by itself establish incapacity and capacity must be assessed for the specific decision with appropriate communication support, an early baseline lets you tell later change from long-standing deficit and lets the difficult conversations happen while the patient can still fully participate.

26.4 The genetics and the C9ORF72 behavioural phenotype

Most disease is sporadic, but a meaningful minority is inherited, and the genes involved are where the motor and psychiatric ends of the spectrum are physically joined. About 5 to 20 percent of cases are familial, and several genes have been implicated, among them SOD1, C9ORF72, VCP, TARDBP, FUS, CHMP2B and ATXN2. One of these deserves separate teaching because it is the gene that most concerns a psychiatrist. Carriers of the **C9ORF72** mutation, an expanded hexanucleotide repeat, a six-base repeat rather than the trinucleotide repeat expansions that underlie Friedreich ataxia and several of the spinocerebellar ataxias described elsewhere in this chapter, may present with behavioural syndromes, especially psychosis and agitation. The same repeat sits on both the motor disease and the frontotemporal dementia, which is precisely why a single gene can produce weakness in one relative and a psychiatric presentation in another.

The clinical consequence is a differential trap in reverse. A middle-aged adult who presents with new psychosis, agitation or a change in social conduct, and who on careful examination has subtle wasting or fasciculations, may be showing the behavioural face of a C9ORF72-related disease rather than a primary psychiatric illness. The primary nosology and management of psychosis itself belong to the Schizophrenia: management, antipsychotics and adverse effects notes; the point owned here is only that this disease can announce itself through the mind before the limbs, and that the family history and the motor examination are the clues that most help separate the

two, though neither, when absent, rules out a C9ORF72-related disease. The corpus does not state a numeric repeat-size threshold for pathogenicity, and none is asserted here.

26.5 Pseudobulbar affect: recognition and its traps

The most distinctive affective disturbance of the disease is **pseudobulbar affect**, a syndrome of disinhibited affective display in which the patient produces inappropriate and exaggerated laughing or weeping in response to environmental stimuli, and it may affect **up to 50 percent** of patients. Its two discriminating features are what an examiner will press. First, the behaviour is not usually associated with any disturbance of the underlying mood, so the weeping patient is not necessarily sad and the laughing patient is not elated. Second, the display tends to be stereotyped for each patient, taking much the same form each time and often disproportionate to or independent of the apparent trigger. The full phenomenology, pathophysiology and management of pseudobulbar affect under any name belong to the Stroke, TBI, delirium and focal brain syndromes notes; what this chapter owns is that it occurs in this disease and must be distinguished from a mood disorder before either is treated.

TRAP **Reading the affective display as the mood.** A patient weeping in the clinic is assumed to be depressed and one laughing to be euphoric, and an antidepressant or a mood stabiliser follows. The defining clue is that in pseudobulbar affect the display is dissociated from the underlying mood and is stereotyped, so the question to ask is not what the tears look like but whether the patient reports the corresponding feeling. Get this wrong and you treat a phantom mood while missing a treatable disorder of affective control.

PITFALL

Confusing two terms that share a prefix. **Pseudobulbar palsy** is the motor sign of upper motor neuron dysfunction in the corticobulbar tracts supplying the bulbar muscles, one of the physical findings listed above, whereas pseudobulbar affect is the emotional syndrome of pathological laughing and crying. They can coexist in this disease and are easily elided, but one is a disorder of the tongue and swallow and the other a disorder of affective display, and only the second answers to the treatments in the next section.

26.6 Treating pseudobulbar affect

Recognising pseudobulbar affect is worthwhile because, unlike much in this disease, it can be treated, and treated on its own indication rather than as though it were the mood disorder it mimics. Successful treatment has been reported with tricyclic compounds or with selective serotonin reuptake inhibitors, used off-label, which makes a low-threshold trial reasonable once the diagnosis is secure. The one agent developed specifically for the syndrome is a combination of dextromethorphan and quinidine, marketed as Nuedexta, which is effective in many cases and is the only agent approved by the United States Food and Drug Administration for pseudobulbar affect. The reference source states these as treatment facts and quotes no milligram doses for them, so none is given here; the fuller pharmacology of the syndrome sits with the Stroke, TBI, delirium and focal brain syndromes notes. For the psychiatrist the practical teaching is compact: name the syndrome correctly, separate it from mood, and reach first for a serotonergic agent or

the approved combination rather than escalating an antidepressant against a depression that is not there.

26.7 Disease-modifying and symptomatic drug treatment

Two established drugs modify the disease itself, and both matter to the psychiatrist for reasons that go beyond survival. The first is **riluzole**, which blocks glutamate release and prolongs survival by about **3 months** in clinical studies. Two details earn their place in a psychiatry chapter. One is that an oral suspension, Tiglutik, exists to avoid the problems of crushing tablets, which is directly relevant in a disease whose bulbar involvement makes swallowing tablets unsafe. The other is a liaison caution: riluzole in overdose may cause agitation, so a disturbed patient on this drug is one in whom the agitation may be pharmacological rather than psychiatric.

The second disease-modifying agent is **edaravone**. Given for amyotrophic lateral sclerosis, in its intravenous formulation it is infused at 60 mg on an initial cycle of daily dosing for 14 days followed by a 14-day drug-free period, and thereafter in cycles of daily dosing on 10 of every 14 days with 14-day drug-free periods between them; an oral suspension formulation also exists. Its benefit in the trials was read on the **ALS Functional Rating Scale-Revised, a twelve-item rating scale spanning fine motor, gross motor, bulbar and respiratory function**, the instrument the reader will meet as the standard functional endpoint in this disease. Neither drug is curative, and both are best understood by the psychiatrist as tools that buy time and function rather than reverse the illness, which keeps the therapeutic conversation honest.

26.8 Prognosis and the end-of-life frame

Every psychiatric decision in this disease is made against a foreshortened prognosis, and stating it plainly is part of the clinician's duty. Death, from respiratory failure or aspiration, typically occurs within 3 years of diagnosis. The one course-modifying decision the source records is the patient's own: the course may occasionally be prolonged, particularly where the patient decides to accept respiratory support. That single fact is what makes advance discussion of ventilation unavoidable rather than optional, and it must be had while cognition allows, because the same disease that shortens survival can, through its frontal involvement, impair the very capacity a valid decision requires.

The wider end-of-life psychiatric care of these patients, the systematic attention to mood and to the psychological weight of a foreseeable death, the formal assessment of decisional capacity as the illness advances, and advance care planning for ventilation and for feeding, is real and important, but it rests on general palliative and psychiatric principles and on the primary management of depressive illness set out in the Mood disorders: pharmacotherapy notes rather than on any figure specific to this disease. Where the numbers are not there to be quoted, the honest teaching is to name the task and treat the patient, not to invent a prevalence. The psychiatrist's contribution across these final years is to keep the mind in view in a disease everyone else is watching for its muscles.

Motor neuron disease is not a pure disorder of movement: in amyotrophic lateral sclerosis, one-third to one-half show frontal-executive cognitive change on the frontotemporal degeneration spectrum, and up to half develop a pseudobulbar affect that is dissociated from mood; cognitive impairment is associated with shorter survival in a disease whose own course is short, with death typically within 3 years of diagnosis.

27 · The functional neurological disorder spectrum and its positive signs

A psychiatrist is often the clinician who recognises this disorder and helps demonstrate the positive signs that establish it. Functional neurological symptom disorder is the modern name for what was long called conversion disorder: genuine neurological symptoms - limb weakness, tremor, seizure-like attacks, visual loss, gait failure - that are real and disabling but arise from a disorder of nervous-system functioning rather than from structural damage to the nervous system. It earns its marks because the whole diagnostic logic has been rewritten in the current classifications, and because the signs that establish it are elegant, reproducible at the bedside, and perfectly suited to a viva. The old teaching treated the disorder as a diagnosis of exclusion, reached once investigations came back clear; the current teaching treats it as a positive, rule-in diagnosis established by specific clinical signs. This section works through those positive signs across the spectrum, then closes on how the diagnosis is explained to the patient and treated.

27.1 A rule-in diagnosis, not a diagnosis of exclusion

The single most examined point in this topic is a change in how the diagnosis is reached. Under DSM-5-TR, **functional neurological symptom disorder**, historically termed conversion disorder, rests on clinical findings that show clear evidence of incompatibility between the symptom and recognised neurological disease. The diagnosis is explicitly not one of exclusion. It is not made simply because investigations are normal, and it is not made because a symptom looks bizarre or defies an anatomical explanation on first impression. What establishes it is positive evidence, a demonstration that the nervous system remains capable of the movement or the perception that the patient cannot voluntarily produce.

The idea at the heart of most positive signs is **internal inconsistency**: a physical sign elicited by one method of examination is no longer present when the same function is tested a different way, or when the patient's attention is drawn elsewhere. Positive signs demonstrate incompatibility through inconsistency or incongruence with recognised neurological disease, and internal inconsistency is one important route rather than a property of every sign; that idea generates much of the bedside repertoire. It also carries a clinically liberating corollary. Because the diagnosis is made on positive signs rather than on the absence of disease, it can be made even in a patient who also has a structural neurological disorder such as epilepsy or multiple sclerosis, according to DSM-5-TR. Functional and organic symptoms coexist in one patient far more often than the older either/or framing allowed. The manual adds a single safeguard against over-reading any one test: the diagnosis should rest on the overall clinical picture and not on a solitary finding.

RULE **Functional neurological disorder is a positive, rule-in diagnosis demonstrated by internal inconsistency, and it can be made alongside genuine neurological disease.** The examiner is testing whether you have absorbed the shift from the exclusion model to the positive-sign model. The presence of a real neurological illness does not exclude a functional overlay, and normal investigations do not establish one.

TRAP **Labelling a symptom functional because the scan is normal or because the presentation looks bizarre.** DSM-5-TR states plainly that the diagnosis is not made on normal investigations and not made because a symptom is bizarre. A normal image does not establish a functional disorder and does not by itself exclude every organic one, and strange-looking symptoms occur in organic disease too. The call is made on a demonstrated incompatibility sign, never on the strangeness of the story or the emptiness of the report.

MNEMONIC

Present one way, absent the other

Every functional positive sign is a demonstration of **internal inconsistency**: the deficit is present when the function is tested directly and absent when the very same function is recruited automatically or through distraction. If you can show the movement returning by another route, you have your rule-in evidence.

27.2 Functional limb weakness: Hoover's sign and the hip abductor sign

Functional limb weakness is the commonest presentation of the disorder and the one with the most reliable signs. The best known is **Hoover's sign**, described by Kaufman for suspected psychogenic weakness of one leg. It exploits the reflex that hip extension in one leg is coupled to hip flexion in the other. When a patient with genuine paresis tries to raise the sound leg, the affected leg presses down; the functional patient, by contrast, fails to press down with the sound leg when raising the weak one, yet unconsciously presses down with the weak leg when attempting to raise the sound one. Elicited formally, weakness of hip extension in the affected leg returns to normal strength when the opposite hip is flexed against resistance, according to DSM-5-TR.

The accompanying figure sets these positive signs out across the spectrum, so that the reader can see how each one turns a claimed deficit into a demonstrable inconsistency.

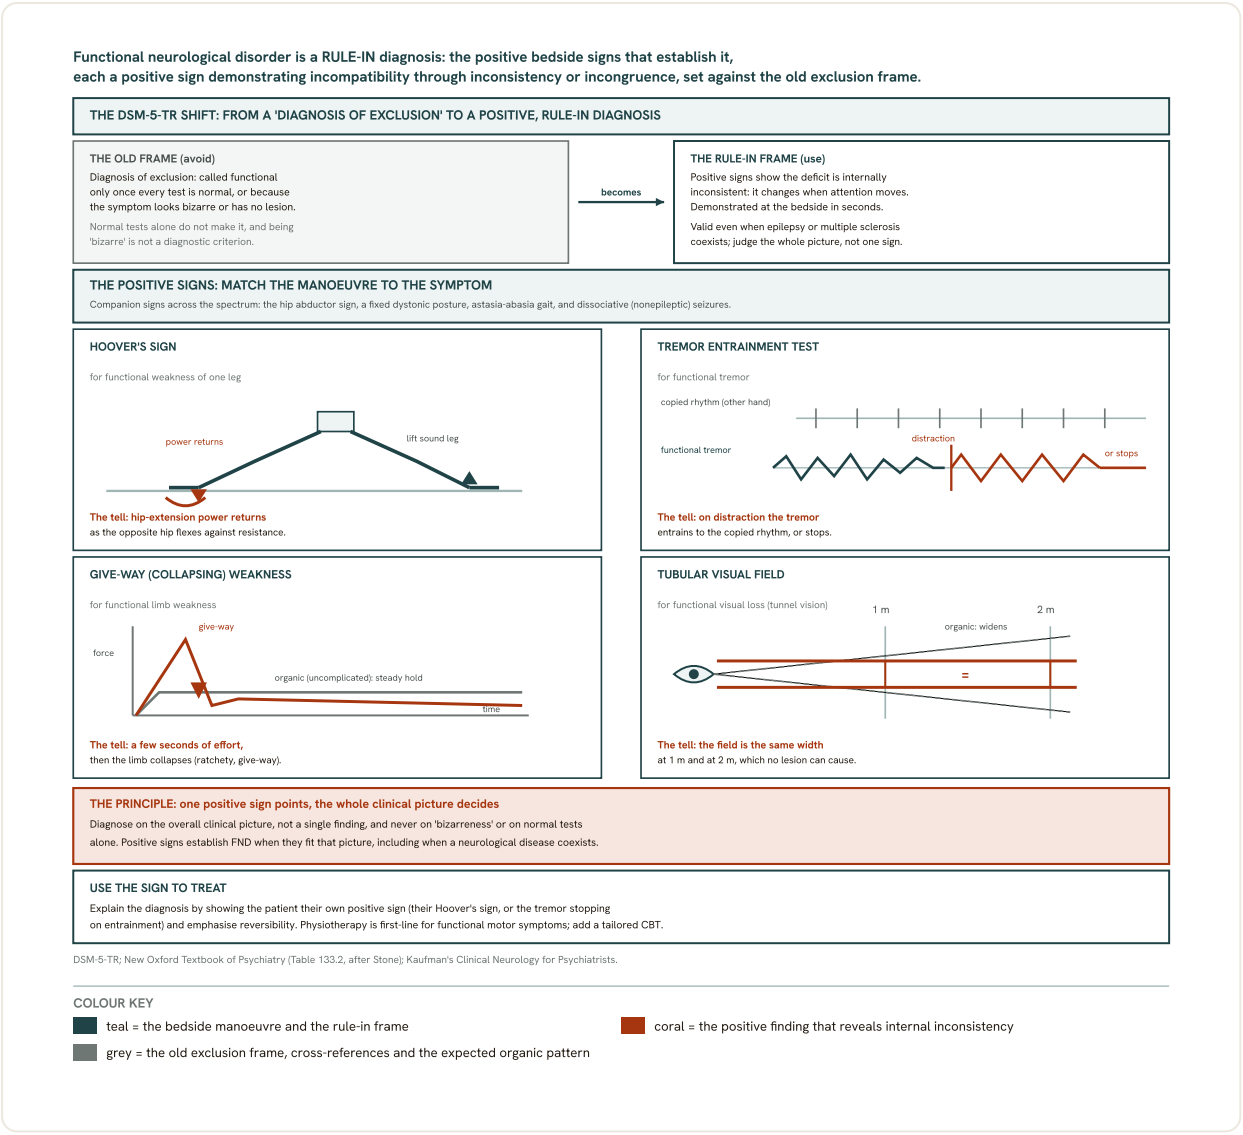


Fig 35.10 **Functional neurological disorder: the positive signs that rule it in.** DSM-5-TR reframes functional neurological symptom disorder as a rule-in diagnosis demonstrated by internal inconsistency, not a diagnosis of exclusion, and it can be made even when a neurological disease such as epilepsy or multiple sclerosis coexists. Four bedside manoeuvres are shown: Hoover's sign (hip-extension power returns as the opposite hip flexes against resistance), the tremor-entrainment test (the tremor locks to a copied rhythm or stops on distraction), give-way or collapsing weakness (a few seconds of effort then collapse), and a tubular visual field (the same width at 1 m and at 2 m). Base the diagnosis on the whole clinical picture and not a single sign, then use the sign to explain the diagnosis, with physiotherapy first-line for functional motor symptoms and a tailored CBT. (DSM-5-TR; New Oxford Textbook of Psychiatry, Table 133.2; Kaufman's Clinical Neurology for Psychiatrists.)

Its companion is the **hip abductor sign**, in which weakness of thigh abduction returns to normal when the opposite hip is abducted against resistance, on the same reflex logic that Hoover exploits for extension. The related discrepancy worth eliciting is a mismatch between two tasks that use the same muscles: weakness of ankle plantar flexion tested on the couch set against a preserved ability to walk on tiptoes, per DSM-5-TR. The value of these manoeuvres is that they compare a voluntary movement against an automatic one and let the automatic movement betray the intact pathway, though they still require adequate comprehension and contralateral effort and must be read alongside pain and cognition rather than as tests of the patient's honesty. That is why they are trusted where an assessment of effort alone would not be.

27.3 Give-way weakness, the face-hand test and the global pattern

Beyond the reflex signs, several patterns in the way weakness behaves point to a functional origin. Kaufman describes the **give-way** or collapsing effort, in which the patient produces a brief exertion of several seconds before the limb returns to an apparently paretic position; sustained organic weakness holds more steadily, whereas this intermittent, ratchety collapse, once pain and fatigue are accounted for, supports a fluctuating and probably psychogenic deficit that is weighed with the other signs rather than read alone. The paired bedside test is the **face-hand test**: the supposedly paralysed hand is held over the patient's own face and released, and a functional patient momentarily exerts enough strength to deflect the falling hand so that it does not strike the face, a preserved protective response that supports, rather than proves, a functional origin.

A further clue is the absence of functional impairment despite the appearance of profound weakness, such as the patient who walks on heels and then on toes even though manual testing seemed to show marked ankle weakness, again per Kaufman. The New Oxford Textbook of Psychiatry adds a distributional sign: weakness that is global, affecting extensors and flexors of a limb equally, rather than following the pyramidal pattern in which extensors of the arm and flexors of the leg give way first. A global, whole-limb pattern can support a functional origin when it is incongruent with the neurological examination, but organic disease can also weaken a limb globally, so the pattern is a pointer to weigh rather than proof on its own.

> **PITFALL**
>
> The commonest clinical error here is to make the call on impression rather than on neuroanatomy. A patient who seems anxious, whose story is inconsistent, or whose weakness feels non-organic to the examining hand has not thereby been shown to have a functional disorder. DSM-5-TR is explicit that the diagnosis is not made because a symptom is bizarre, and that it rests on the overall picture rather than a single impression. Elicit an incompatibility sign, or reserve judgement; a psychogenic label pinned on a gestalt is how genuine disease gets missed.

27.4 Functional tremor and functional dystonia

Functional movement disorders carry their own rule-in tests, and the most powerful is the **tremor entrainment test** for functional tremor. The patient is asked to copy the examiner in making a rhythmical movement with the contralateral hand or foot, which loads the same motor-timing resource the tremor is drawing on. The test is positive, per DSM-5-TR, when the tremor takes up the imposed rhythm, when it is suppressed altogether, or when the patient simply cannot copy a simple rhythmical movement at all. Additional supportive features of a functional limb tremor are variability in its frequency or in its direction, neither of which is characteristic of the organic tremors it mimics, whose frequency is comparatively fixed.

The other cardinal functional movement disorder is **functional dystonia**. Its typical presentation, per DSM-5-TR, is a fixed inverted position of the ankle, a clenched fist, or a unilateral contraction of the platysma, characteristically with sudden onset. The New Oxford Textbook of Psychiatry describes the same picture as a fixed dystonic posture, most often of the

hand with flexion of the fingers, wrist and elbow, or of the ankle with plantar flexion and inversion. The abrupt onset and the fixed, non-mobile quality of the posture distinguish it from most organic dystonias, which tend to be mobile, task-specific and gradual in evolution. Fatigable or fluctuating weakness of a limb should also raise the neuromuscular-junction and dystrophic disorders that enter this differential, which are taken up elsewhere in these notes rather than here.

27.5 Functional (dissociative) seizures and the duration question

Attacks that resemble epileptic seizures or syncope are among the most important and most litigated presentations of the disorder. They carry several names for the same thing, **functional or dissociative seizures**, also called non-epileptic seizures or attacks, and the terminology itself signals that they are not driven by epileptic discharge. The positive features that suggest a functional origin, according to DSM-5-TR, include persistent eye closure during the attack, sometimes with active resistance to having the eyes opened, bilateral motor movements occurring with preserved awareness, and a duration of the event longer than the norm for an epileptic seizure. The New Oxford Textbook of Psychiatry adds ictal weeping, crying during or immediately after the attack, the ability afterwards to recall the experience of being in a generalised shaking attack, and ictal hyperventilation, since respiration ceases in a generalised epileptic seizure but commonly speeds up during a non-epileptic one.

These features must be combined rather than used singly, and they may be supported by a normal simultaneous ictal electroencephalogram, though DSM-5-TR cautions that a normal ictal recording alone does not exclude every form of epilepsy or of syncope. Duration is the classic quantified clue, and it is also the classic source of examination error.

TRAP **Treating a single attack-duration cut-off as diagnostic, and assuming the standard references agree on it.** DSM-5-TR gives the suggestive feature as a duration **longer than 5 minutes**. The New Oxford Textbook of Psychiatry sets the long-duration threshold appreciably shorter, so the two authorities a candidate is likely to have read give different numbers for the same construct. Both stress that duration is only supportive and must be combined with other features. The safe answer names the disagreement and refuses to let any single duration stand as proof.

IN INDIA

Functional and dissociative attacks can coexist with epilepsy, and where epilepsy and endemic seizure-causing infections such as neurocysticercosis are common, the pull toward settling every atypical attack as one or the other is strong. The clinical decision is to diagnose the functional attack on its own positive signs, ideally captured on simultaneous electroencephalography, and never to assume that finding a genuine seizure disorder excludes a functional one, nor that a functional diagnosis excludes coexisting epilepsy. Both can be present, and each is treated on its own terms.

27.6 Functional gait and the value of distraction

Gait is where a functional disorder is often most visible, and its signature is a mismatch between apparent danger and actual safety. The most readily identifiable pattern, per Kaufman, is **astasia-abasia**, from the Greek for an inability to stand and an inability to walk. The patient staggers, balances momentarily and appears to be in great danger of falling, yet catches themselves at the last moment by grabbing railings, furniture and even the examiner, and never actually injures themselves. That preserved postural control under apparent chaos, the fall that is always averted, is the giveaway, because it demonstrates an intact balance system working automatically beneath a gait that is failing on demand. The New Oxford Textbook of Psychiatry describes a related dragging gait in which the forefoot stays in contact with the ground and the hip is held externally or internally rotated, a pattern that supports a functional gait when it is inconsistent or incongruent with the examination, though organic disease can drag the forefoot too.

Distraction is the examiner's most useful single tool across this spectrum, and it is nowhere more telling than in gait and stance. In the New Oxford description, a patient with an apparently positive Romberg's test is asked to guess numbers traced on their back or to perform a complex motor task, and the sway settles; in a functional gait problem, balance improves significantly once attention is diverted from the act of standing. The principle is the same one that runs through every positive sign: engage automatic motor control, take conscious attention off the failing function, and the function returns. Distraction does not trick the patient so much as reveal the pathway that voluntary effort could not access.

27.7 Functional visual and speech symptoms

The sensory and speech domains complete the spectrum, and they too are diagnosed on internal inconsistency rather than on a normal structural examination. In functional speech, the positive sign is an internal inconsistency in articulation and phonation, per DSM-5-TR: a whispering dysphonia or a stutter that shifts or resolves when the patient is off guard, coughs normally, or is engaged in unselfconscious speech. In the visual domain, the hallmark is a **tubular visual field**, or tunnel vision, together with tests that show internal inconsistency in visual acuity. The New Oxford Textbook of Psychiatry gives the operational definition that exposes it: a tubular field has the same width at **1 m as at 2 m**, which is geometrically impossible for a genuine field defect, since a real constricted field widens as the target moves further from the eye.

Acuity is probed with the fogging test, per DSM-5-TR, in which the patient reads a chart with both eyes open while the good eye is progressively fogged with lenses, so that any letters still read must be the work of the supposedly blind eye. The table below consolidates the positive signs across the whole spectrum, since the examinable core is not any one sign but the shared logic that binds them.

DOMAIN	POSITIVE RULE-IN SIGN	WHY IT IS INCOMPATIBLE WITH ORGANIC DISEASE
Limb weakness	Hoover's sign; the hip abductor sign	Power returns via the contralateral reflex the patient cannot suppress
Weakness pattern	Give-way effort; face-hand test; global weakness	Collapsing, whole-limb weakness ignores the pyramidal distribution
Tremor	Tremor entrainment; variable frequency	Tremor entrains or stops when the timing resource is loaded
Dystonia	Fixed inverted ankle or clenched fist, sudden onset	Fixed, abrupt posture unlike the mobile organic dystonias
Seizures	Eye closure with resistance; preserved awareness	Retained awareness during bilateral shaking; a normal ictal EEG supports but does not prove it
Gait	Astasia-abasia; improvement on distraction	Balance is intact automatically but fails on demand
Vision	Tubular field; fogging test	A field that does not widen with distance cannot be organic

27.8 Explaining and treating the disorder

The first therapeutic act is the diagnosis itself. Management, according to the New Oxford Textbook of Psychiatry, should usually be multi-disciplinary, and its foundation is a plausible and agreed explanation of the diagnosis followed by effective triage to the right therapy. The explanation is not a reassurance script. It has recognisable components, and they are what a viva answer on this topic is expected to name.

- Take the problem seriously.
- Give the patient a positive diagnostic label rather than a list of the things they do not have.
- Demonstrate the rationale for the diagnosis, by showing them their own positive Hoover's sign or how their tremor stops during the entrainment test.
- Emphasise potential reversibility, often through the metaphor that the trouble lies in the running of the system rather than in its structure.
- Provide written information the patient can take away.

GOOD TO REMEMBER

Showing the patient the sign is both the proof of the diagnosis and the start of the treatment. When a patient watches their own leg regain power as the opposite hip flexes, or their tremor dissolve the moment they tap a rhythm with the other hand, the demonstration does what no reassurance can: it makes the reversibility credible from the inside. Build the explanation around the sign you elicited, not around a scan that was normal.

On the specific therapies, the New Oxford Textbook of Psychiatry holds that physiotherapy should be a first-line treatment for functional movement disorders and limb weakness, provided the patient has some confidence in the diagnosis and is motivated to change. The physiotherapy has to be the tailored, functional-disorder kind; generic physiotherapy of the sort delivered for stroke may in fact be counterproductive here, because it reinforces the very abnormal patterns the treatment is trying to unlearn. Alongside it, a tailored cognitive behavioural therapy delivered as an informed guided self-help programme was found to be better than treatment as usual, and a specific form of cognitive behavioural therapy has been developed for dissociative seizures. The three pillars to carry are therefore a credible explanation, targeted physiotherapy and a tailored psychological therapy.

over 5 min	1 m = 2 m
ATTACK DURATION THAT FLAGS A FUNCTIONAL SEIZURE (DSM-5-TR)	TUBULAR FIELD: EQUAL WIDTH NEAR AND FAR

Functional neurological disorder is diagnosed by positive signs of internal inconsistency, not by normal scans, and the sign you demonstrate to the patient is the first step of the treatment.

28 · Myasthenia gravis, mitochondrial and muscular dystrophy neuropsychiatry

Three very different diseases meet here for one reason: each disturbs muscle, nerve terminal or mitochondrion, yet each reaches the psychiatric clinic wearing a mental disguise. Myasthenia gravis produces a fluctuating fatigable weakness that is repeatedly mistaken for a functional disorder or for depressive exhaustion; the mitochondrial disorder MELAS can open with a psychosis years before its neurology declares itself; and myotonic dystrophy carries an almost inescapable cognitive decline and a flattening of personality. The marks in this topic sit at two junctions the examiner returns to: the point where an organic neuromuscular illness is misread as psychiatric, and the point where a psychotropic drug becomes hazardous in a person whose muscles, heart or metabolism are already compromised. This account works through myasthenia first, then the mitochondrial and dystrophic disorders, keeping to what the primary neurology and prescribing sources state and flagging, at each step, the confidently remembered number that is not safe to assert.

28.1 The neuromuscular-junction lesion, and the mind it leaves untouched

Myasthenia gravis is the classic disorder of the neuromuscular junction. Circulating antibodies to the **nicotinic acetylcholine receptor** block the binding of acetylcholine on the postsynaptic membrane, so repeated stimulation fails to sustain contraction and the muscle fatigues. The distribution is characteristic: the extraocular, facial, neck and proximal limb muscles bear the brunt, and it is weakness spreading to the respiratory muscles that threatens life. The single fact a psychiatrist must hold above all others is where these antibodies do not act. They attack the muscle-type nicotinic receptors of the junction, they do not cross the blood-brain barrier, and they ordinarily do not reach or attack central cholinergic receptors, so the disease itself does not

directly disturb consciousness or cognition. Extraocular muscles weaken while the intraocular pupillary muscles stay strong, and bladder, bowel and sensory function are all preserved. The antibodies do, however, cross the placenta, which is why an affected mother can transmit a transient neonatal myasthenia to her infant. This split, a peripheral postsynaptic lesion with an intact central nervous system, is the whole reason myasthenia belongs in a psychiatric account: it is dramatic weakness without a single central sign.

RULE **Myasthenia gravis weakens the body and spares the mind.** Profound, even ventilator-dependent weakness coexists with an entirely normal mentation and level of consciousness, because the antibody is confined to the peripheral nicotinic junction. A patient who is weak and also drowsy, confused or cognitively altered has something else, or something in addition, going on. That dissociation between striking weakness and a clear sensorium is the feature that separates myasthenia from both a primary psychiatric disorder and a central neurological one.

28.2 The fatigable, diurnal signature and how the diagnosis is secured

The hallmark of myasthenia is **fatigable weakness**: strength that is present at rest and drains with use, so power is worst after sustained effort and recovers with rest and sleep. The pattern is almost diurnal, mild in the morning and heaviest in the late afternoon and evening, and it fluctuates from one day to the next. The eyes usually announce the disease. Almost **90%** of patients, typically young women or older men, present first with diplopia and ptosis. Facial and neck involvement follows, giving a nasal, dysarthric speech and a snarling grimace on attempting to smile, and progression to the neck, shoulder, swallowing and respiratory muscles produces a bulbar palsy and, at its most severe, respiratory distress. The important clinical texture is that the weakness is asymmetric and variable rather than fixed, which is exactly why it is so easily doubted.

Diagnosis rests on the bedside and the laboratory rather than on clinical impression, and several tests are used together:

- The **Tensilon (edrophonium) test** and the ice-pack test, both of which briefly reverse a ptosis.
- Serum antibody assays, for acetylcholine-receptor and, in seronegative cases, muscle-specific-kinase antibodies.
- Repetitive nerve stimulation showing a decremental response as the junction fatigues.

28.3 Serology, the thymus and the treatment ladder

The majority of patients are seropositive for acetylcholine-receptor antibodies; of the seronegative remainder, about half carry antibodies to **muscle-specific kinase (MuSK)**, and a smaller group are seronegative on current assays. The disease also keeps company with the thymus and the thyroid, and both associations are worth carrying because treating them can improve the myasthenia itself.

~80%	10%	5%
SEROPOSITIVE FOR ACHR ANTIBODIES	HAVE A MEDIASTINAL THYMOMA	HAVE UNDERLYING HYPERTHYROIDISM

Treatment follows two logics that are worth keeping distinct. The first raises acetylcholine at the junction: the cholinesterase inhibitor pyridostigmine is the mainstay, with parenteral neostigmine when a patient cannot swallow. The second restores the receptor by damping the immune attack, using corticosteroids, other immunosuppressants, plasmapheresis or intravenous immunoglobulin, and, in refractory disease, the anti-CD20 monoclonal antibody rituximab. The division matters to the psychiatrist because it is the immunomodulatory arm, and corticosteroids above all, that carries the psychiatric hazard taken up below.

28.4 The organic illness most often mistaken for a functional one

This is the reason myasthenia earns a place in a psychiatry syllabus at all. A patient may present with the abrupt onset of dysarthria, dysphagia and diplopia and be labelled, wrongly, with a conversion disorder. Everything about the surface presentation invites the error. The weakness fluctuates and is relieved by rest, so it looks inconsistent; the objective signs are often transient or subtle, as they are in multiple sclerosis, focal seizures and small strokes, so a purely history-based assessment can mislead; exacerbations usually arise spontaneously but can be precipitated by an intercurrent illness or by psychological stress, which invites a psychogenic reading; and, because mentation is untouched, nothing in the mental state points away from a primary psychiatric cause. The functional neurological disorder spectrum and its positive signs are taken up separately in these notes; the point to hold here is the discipline of excluding myasthenia before that label is applied.

TRAP **Reading fatigable, stress-sensitive, evening-worse weakness as conversion or as depressive fatigue.** Fluctuation, relief by rest and provocation by stress are typical of myasthenia, not evidence against it, and a normal sensorium removes the one reassurance a clinician might reach for. The safe rule is that new bulbar or ocular fatigable weakness earns antibody testing and repetitive nerve stimulation before it earns a functional or a depressive label, because the price of missing myasthenia is a preventable respiratory crisis.

28.5 Prescribing in myasthenia: steroids, sedation and neuromuscular safety

Two prescribing problems converge in myasthenia. The first is that its own immunotherapy is psychoactive. Corticosteroids, a mainstay of treatment, cause depression, anxiety, psychosis and mania; short courses tend toward euphoria or mania and long courses toward depression; onset is usually rapid, within one to two weeks and by one review a median of about **11.5 days**; and the incidence is dose-related. Compared with others who have the same underlying illness, patients on glucocorticoids are almost twice as likely to become depressed, four times as likely to become manic and five times as likely to become delirious, and withdrawal, especially if rapid, can itself precipitate symptoms. The dose relationship, from the Boston Collaborative Drug Surveillance Program, is worth carrying as a table.

PREDNISONE DOSE	ACUTE PSYCHIATRIC REACTION
Under 40 mg daily	1.3%
40 to 80 mg daily	4.6%
More than 80 mg daily	18%

The second problem is that several psychotropics are unsafe in a disease of respiratory-muscle weakness. Because the bulbar and respiratory muscles can fail, drugs that depress respiration are hazardous, and benzodiazepines and the Z-drug hypnotics such as zopiclone are listed as contraindicated in myasthenia gravis for exactly that reason. Other agents worsen neuromuscular transmission directly: propranolol may aggravate myasthenic weakness, and the anticholinergic action of quinidine can do the same. The general choice and adverse-effect profile of antipsychotics and hypnotics belongs to the Schizophrenia: management, antipsychotics and adverse effects notes; what this topic adds is that the neuromuscular disease reorders those safety priorities.

PITFALL

Reaching for a benzodiazepine or a Z-drug hypnotic to settle anxiety or insomnia in a known myasthenic. In a disease whose danger is respiratory-muscle failure, a respiratory depressant is precisely the wrong sedative, and the caution holds even when the request seems minor. Where sedation or anxiolysis is genuinely needed, it is the neuromuscular risk, not habit, that should drive the choice of agent.

28.6 MELAS: the commonest mitochondrial disorder and its genetics

MELAS, mitochondrial encephalomyopathy with lactic acidosis and stroke-like episodes, is the most common of the mitochondrial disorders and the one most likely to surface in a psychiatric differential. Its genetics are maternal, because mitochondria are inherited through the egg, and about **80%** of cases arise from a single A3243G point mutation in the tRNA-Leu(UUR) gene; a further tenth involve other regions of that same gene and the remainder involve other mitochondrial genes. Maternal transmission also helps explain why severity and family history are so variable, since the mutant load carried by different tissues need not be equal.

The clinical core is in the name. The **stroke-like episodes** are the defining event, and their crucial feature is that the lesions do not conform to arterial territories and may involve grey or subcortical white matter, which is what marks them out from a true ischaemic stroke. Onset is usually in childhood, after a period of normal development, though a first presentation in early adulthood is well described. The syndrome is multisystem:

- Short stature, migraine, recurrent vomiting and seizures.
- An elevated lactate, whose level has been correlated with the burden of neurological symptoms.
- Over time, deafness, diabetes, retinopathy and cardiac abnormalities, in a course that ranges from a single episode to a relentless accumulation of strokes.

28.7 Where the lesions fall, the psychosis that comes first, and treatment

The stroke-like lesions favour the posterior brain. They fall in the temporo-parieto-occipital cortex, and also in the basal ganglia, brainstem and cerebellum, and the posterior, parieto-occipital emphasis explains the hemianopia and cortical blindness that punctuate the illness, while hemiparesis follows involvement of the motor cortex and descending motor pathways. For

the psychiatrist the more treacherous fact is what can come first. MELAS has been reported to present with a schizophrenia-like psychosis, often in the third decade, several years to a decade before the neurological diagnosis is made, and cases resembling depression, obsessive-compulsive disorder, bipolar disorder and borderline personality disorder are described. In each the course runs atypically, and a retrospective look frequently turns up previously unremarked short stature, diabetes or unexplained somatic complaints; it is usually the later neurological signs, cognitive decline and strokes on imaging that finally reveal the diagnosis.

TRAP **Taking an atypical early-adult psychosis at face value when it sits on a mitochondrial background.** A psychosis in the third decade with an unusual course, set against short stature, sensorineural deafness, diabetes or a maternal line of similar illness, should prompt thought of MELAS rather than a purely primary psychosis. The mental picture can precede the strokes by years, so it is the somatic clues, not the mental state alone, that open the door to the diagnosis.

There is no proven disease-modifying treatment. A study of 13 patients given L-arginine during acute stroke-like episodes reported a reduction in their frequency and severity without altering the overall trajectory, and care is otherwise symptomatic: cochlear implants for deafness, anticonvulsants for seizures, antipsychotics for psychosis, and treatment of the diabetes.

GOOD TO REMEMBER

Because MELAS itself predisposes to diabetes, an antipsychotic with a high risk of glucose intolerance, such as olanzapine, is best avoided or used with caution when psychosis complicates the disorder. The reasoning generalises: in any disease that already loads the metabolic dice, the metabolic profile of the antipsychotic should weigh more heavily than usual in the choice.

28.8 Myotonic dystrophy: an autosomal dominant channelopathy

Myotonic dystrophy is the most frequent myopathy of adult life, usually declaring itself between 20 and 25 years and affecting both sexes equally. Its lesion is an excessive CTG trinucleotide repeat in the DMPK gene on chromosome 19. Because the expansion disturbs ion channels in the membranes of muscle and other cells, the disease is classed as a **channelopathy**; it is transmitted as an autosomal dominant trait and, like other unstable-repeat disorders, shows **anticipation**, appearing earlier and more severely down the generations. The distinct type 2 disease, proximal myotonic myopathy, instead carries a four-nucleotide repeat, shows no genotype-phenotype correlation and does not anticipate.

The clinical signature is named for its cardinal sign. **Myotonia** is an involuntary, prolonged muscle contraction that delays release of the grip for several seconds after a handshake, and percussion of the thenar eminence produces a visible sustained contraction. Weakness and wasting are distal and facial, and atrophy of the face and temples gives the characteristic **hatchet face** with ptosis and a prominent forehead. The disorder is multisystem, with cataracts, cardiac conduction disturbance and endocrine failure, testicular atrophy, diabetes and infertility, running alongside the muscle disease. This distal, myotonic picture is the mirror image of the proximal weakness and pseudohypertrophy of Duchenne muscular dystrophy.

FEATURE	MYOTONIC DYSTROPHY	DUCHENNE MUSCULAR DYSTROPHY
Weakness distribution	Distal and facial	Proximal
Hallmark muscle sign	Myotonia, grip and percussion	Pseudohypertrophy
Cognitive course	Progressive, increases with age	Non-progressive

MNEMONIC

Grip, Face, Lens, Heart, Gland

The multisystem reach of myotonic dystrophy: a sustained **grip** and percussion myotonia, the hatchet **face**, cataracts in the **lens**, conduction disease in the **heart**, and endocrine **gland** failure. Each element is drawn from the disease as described here and none is added to it.

28.9 Cognition and personality in myotonic dystrophy

The neuropsychiatric heart of myotonic dystrophy is easy to overlook beside the striking muscle signs, yet it is what most shapes the patient's life and the psychiatrist's involvement. Cognitive impairment is almost uniform and, unlike the fixed deficit of Duchenne, it increases with age, so the picture is one of slow decline rather than a stable baseline. Alongside it runs a characteristic change in personality: a lack of initiative and a progressive blandness, an apathetic flattening of drive and affect that families often notice before any formal testing does. The combination, an age-related cognitive decline with apathy and emotional blunting, can be mistaken for a primary depressive disorder or for an independent dementia, and it is better read as part of the multisystem disease than as a separate psychiatric illness. Formal cognitive assessment and support, and attention to the apathy as a source of disability in its own right, are the practical response, and they are what a psychiatrist most often adds to the care of these patients.

In myasthenia gravis the antibody spares the mind while the body fails, so fatigable weakness must be excluded before a functional or depressive label; MELAS can open with a third-decade psychosis years before its stroke-like episodes; and myotonic dystrophy pairs an age-increasing cognitive decline with an apathetic, bland personality.

Consolidate and apply

29 · Worked vignettes

A vignette is where recall becomes clinical judgement. The examiner rarely asks a candidate to define a disease of the brain; far more often a short stem describes a patient, plants one or two features that mark the presentation as something other than a primary psychiatric disorder, and waits to see whether the candidate reads them. This closing section teaches no new disease. It takes the syndromes assembled across this chapter and rehearses them as **recognition signatures**, the compact clinical gestalts that let a reader move from a paragraph of history to the right family of disease under examination time pressure.

The skill being tested is discrimination, not enumeration. Two stems can share a headline symptom, a first psychotic episode or a low mood, and diverge on a single detail that changes the whole answer. Learning the chapter as a set of paired look-alikes, each separated by the one feature that changes the answer, is worth more marks than a longer catalogue of associations, because it mirrors what the clinician actually does at the bedside.

29.1 Read the stem for its planted discriminator

The first task with any case stem is to decide whether the psychiatric picture might be a **secondary neuropsychiatric syndrome**, the mental-state expression of a systemic or neurological disease, rather than a primary disorder in its own right. The chapter's diseases share a small set of tells, and an efficient reader scans for them before committing. Four axes carry most of the discriminating information.

Tempo	Company	Atypia	Background
HOW FAST IT CAME	THE SIGN BESIDE THE SYMPTOM	WRONG FOR A PRIMARY DISORDER	EXPOSURE AND HEREDITY

Tempo asks how quickly the change arrived and whether it fluctuates; an abrupt or subacute course, or a waxing and waning conscious level, is unusual for most primary disorders. Company asks what neurological sign sits beside the psychiatric symptom: a movement disorder, a seizure, a focal deficit or autonomic instability. Atypia asks whether any feature is simply wrong for the diagnosis being entertained, such as a first psychotic episode in later life, visual rather than auditory hallucinations, or failure to respond to adequate standard treatment. Background asks what the patient carries into the illness, an infective or immune exposure, or a movement disorder running through the family tree. The single feature that reorganises the differential is the **discriminator**, and good stems are written around it.

TRAP **Diagnostic overshadowing: attributing new organic symptoms to a psychiatric label the patient already carries.** When a stem opens with a known diagnosis, the tempting move is to read every later feature as part of it, so a new movement disorder becomes a medication effect and a subacute confusion becomes relapse. **Diagnostic overshadowing** is an important way a secondary cause is missed, and examiners plant a pre-existing diagnosis precisely to test whether the candidate keeps the differential open.

29.2 The hyperkinetic case: chorea and its neuropsychiatric tail

A **hyperkinetic** stem leads with excess movement, and chorea is the movement that most often carries a neuropsychiatric tail into the examination. The archetype is a progressive chorea with personality change, irritability and cognitive decline, often with a family history reaching back through affected parents; this pattern points to Huntington disease, whose genetics and psychiatric features are developed in full within this chapter. Against it sits a usually self-limiting chorea in a child or adolescent after a streptococcal illness, sometimes with emotional lability and obsessive features, which points instead to Sydenham chorea. Age, tempo, heredity and whether the course remits separate the two.

Wilson disease enters the same differential from a different door: a young patient with a movement disorder, its characteristic systemic and ophthalmological signs, and psychiatric or behavioural change. Its movement-disorder and neuropsychiatric presentation belongs to this chapter's hyperkinetic overlap, while the hepatic and metabolic disease is developed in the Endocrine and metabolic psychiatry notes. Tics are the other hyperkinetic pattern a stem may use as a distractor; the archetypal tic disorder, Tourette syndrome, and its diagnostic criteria belong to the Child behavioural, emotional and other disorders notes, and recognising a tic as a tic keeps it out of the chorea differential. The recognition point is that a young patient with an involuntary movement disorder and a psychiatric change deserves a search for a defined neurological or metabolic cause before the psychiatric label is fixed.

29.3 The hypokinetic case: parkinsonism with a mental-state change

A **hypokinetic** stem leads with slowness, rigidity and tremor, and the neuropsychiatric question is which parkinsonian syndrome the picture represents and what its mental-state features signify. In established Parkinson disease, the classic examination trigger is well-formed visual hallucinations with initially preserved insight, the presentation of Parkinson disease psychosis, where management turns on antipsychotic selection that does not worsen motor function; that treatment is set out within this chapter's Parkinson disease sections. A slowly accruing dysexecutive and visuospatial decline late in the disease course points instead to Parkinson disease dementia, also developed here.

The discriminator the examiner most often plants is timing and drug response. When cognitive impairment and visual hallucinations arrive around the onset of the motor syndrome rather than well after it, the temporal rule that separates dementia with Lewy bodies from Parkinson disease dementia becomes the issue, and that rule, with the neuroleptic sensitivity that accompanies it, belongs to the Dementias and neurocognitive disorders notes. When parkinsonism responds poorly to dopaminergic therapy, or is joined early by falls, autonomic failure, gaze restriction or asymmetric limb apraxia, the stem is steering away from idiopathic disease toward atypical parkinsonism, which this chapter develops among the degenerative syndromes. Reading a parkinsonian stem is therefore a matter of placing the mental-state change in the sequence of motor, cognitive and autonomic events.

29.4 The subacute encephalopathy: new psychosis with neurological company

Some of the highest-yield stems describe a previously well adult who develops, over a subacute course, a new psychosis or severe behavioural change accompanied by seizures, a movement disorder, autonomic instability or a fluctuating conscious level. This **subacute encephalopathy** pattern is the signature of the autoimmune and infective encephalitides. Anti-NMDA receptor encephalitis is the archetype, often opening with prominent psychiatric features before the neurological picture declares itself; the antibody-mediated limbic encephalitis spectrum and the viral encephalitides, herpes simplex chief among them, produce overlapping pictures and are each developed in full within this chapter.

RULE **One brain, one diagnosis first.** When a new psychiatric syndrome keeps company with a movement disorder, seizures, autonomic instability or a focal neurological sign, the parsimonious reading is a single disease of the brain producing both, not two coincident illnesses. Coincident illness is the less parsimonious reading, weighed in parallel while the unifying organic cause is pursued first.

The examination trap in this territory is that the encephalitides can produce catatonic and psychotic phenomena that mimic a primary disorder. Catatonia phenomenology, its structured rating and the lorazepam challenge belong to the Other psychotic disorders, delusional syndromes and catatonia notes, and the distinction between an antibody-mediated encephalitis and a primary catatonia is developed within this chapter's autoimmune-psychosis discussion. The primary nosology of schizophrenia and of mood disorder, against which these secondary syndromes are read, belongs to that same Schizophrenia notes and to the Mood disorders: pharmacotherapy notes. The recognition point is tempo with company: a psychosis that arrives quickly, alongside a neurological sign, earns investigation for an encephalitis rather than a reflexive primary diagnosis.

29.5 The chronic infection: reading the stem in Indian practice

Endemic exposure rewrites the differential. A stem describing insidious personality change, cognitive decline, mood disturbance or new psychosis on the background of a chronic infection is a recurring examination pattern, and the causes that matter most in Indian practice are not always those a Western textbook lists first. Neurosyphilis, presenting as general paresis of the insane, remains a classic cause of a dementing and disinhibited picture and must stay in the differential of a progressive neuropsychiatric decline. Neurocysticercosis is a leading cause of adult-onset seizures in endemic regions, so seizures with behavioural or psychotic features should raise it. Japanese encephalitis and cerebral malaria produce acute encephalopathic illness and later neuropsychiatric sequelae, and subacute sclerosing panencephalitis presents in a child or adolescent with cognitive and behavioural decline, myoclonus and progressive deterioration. Central nervous system tuberculosis underlies a proportion of subacute meningoencephalitic presentations. Each is developed in full within this chapter.

IN INDIA

Where these infections are endemic, keep them at the front of the differential rather than at the end. In an adult with new-onset seizures and behavioural change, pursue neuroimaging for neurocysticercosis before settling on a primary psychiatric explanation. In a child or adolescent with declining school performance, personality change and myoclonus, hold subacute sclerosing panencephalitis in view. In a progressive dementing syndrome with disinhibition, retain neurosyphilis as a treatable cause worth testing for. The decision that changes management is to investigate the endemic infection early, because several are treatable and the rest still alter supportive care, prognosis and counselling, and each is missed when the presentation is read as psychiatric first.

29.6 The HIV case: an ordinary presentation with an extraordinary differential

A stem that supplies HIV status, or the risk factors for it, invites the candidate to think through the neuropsychiatry of HIV rather than to treat the psychiatric syndrome in isolation. Cognitive slowing, apathy and executive difficulty on this background raise the spectrum of HIV-associated neurocognitive disorders, whose classification and the effect of antiretroviral therapy are developed in full within this chapter. New focal deficits, seizures or a declining conscious level raise the central nervous system opportunistic infections, including the opportunistic viral leukoencephalopathy that this chapter owns among them, progressive multifocal leukoencephalopathy. Depression, mania, psychosis, anxiety and an elevated suicide risk are all recognised psychiatric comorbidities, and the direct neuropsychiatric effects of the antiretroviral agents themselves belong to the chapter's discussion of drug interactions and adverse effects.

GOOD TO REMEMBER

When a psychiatric syndrome and a treatable disease of the brain appear in the same patient, treating one is not an alternative to treating the other. The examination answer, and the clinical one, is **two-front**: correct the underlying disease and manage the psychiatric syndrome in parallel, because a mood or psychotic disturbance driven by an active brain disease seldom settles on psychiatric treatment alone.

Bedside cognitive testing frames the recognition but does not make the diagnosis; the standardised cognitive instruments used to quantify impairment belong to the Psychotherapies, clinical assessment, MSE and nosology notes, and here they serve only to flag that a cognitive syndrome is present and to track its course. The recognition point is that HIV converts an ordinary psychiatric presentation into a differential that must include neurocognitive disorder and opportunistic infection before it is called primary.

29.7 The demyelinating, neuromuscular and functional case

Relapsing neurological events separated in time and space, joined by mood disturbance, cognitive change or pathological affect, describe multiple sclerosis, whose psychiatric manifestations and the effect of disease-modifying therapy are developed in full within this chapter. Multiple sclerosis and motor neuron disease can both produce pseudobulbar affect, the pathological laughing and crying whose full phenomenology, mechanism and treatment belong to the Stroke, TBI, delirium and focal brain syndromes notes; here it is named as a feature these diseases carry. Motor neuron disease itself, through the amyotrophic lateral sclerosis and frontotemporal dementia spectrum, brings behavioural and cognitive change alongside progressive weakness and the end-of-life questions this chapter addresses.

Two neuromuscular tells reward recognition. Weakness that fatigues through the day and recovers with rest, sometimes with the psychiatric consequences of a chronic disabling illness, points toward myasthenia gravis and its neuromuscular-junction cautions, developed here among the neuromuscular syndromes. Weakness, sensory loss or movement that is internally inconsistent and incongruent with neuroanatomy, and that is confirmed by **positive signs** rather

than by exclusion, points toward functional neurological disorder, whose positive-sign diagnosis this chapter develops in full.

PITFALL

Diagnosing functional neurological disorder on impression, because the presentation seems atypical or the patient distressed, rather than on the positive signs that establish it. A functional diagnosis reached by exclusion is unsafe: it both mislabels genuine neurological disease and denies the patient a confident, rule-in diagnosis that carries its own treatment. The call rests on demonstrable incongruence and specific positive signs, not on the clinician's impression.

29.8 Look-alikes and the master recognition spine

Consolidation is easiest as pairs. The table below sets the chapter's most confusable presentations against the single feature that separates them, and the accompanying figure arranges the same discriminators along the recognition axes.

PRESENTATIONS THAT GET CONFUSED	THE DISCRIMINATING FEATURE	POINTS TOWARD
Usually self-limiting chorea after infection in the young against relentless chorea with cognitive decline and family history	Course and heredity	Sydenham chorea against Huntington disease
New psychosis with seizures, a movement disorder or autonomic instability against new psychosis alone	Neurological company and tempo	Antibody-mediated encephalitis against a primary psychotic disorder
Asymmetric parkinsonism responsive to dopaminergic therapy against parkinsonism with early falls, autonomic failure or poor response	Drug response and early red-flag signs	Idiopathic Parkinson disease against atypical parkinsonism
Visual hallucinations and cognitive decline arriving around the onset of parkinsonism against well after established motor disease	The temporal sequence	Dementia with Lewy bodies against Parkinson disease dementia
Fatigable weakness worse through the day against fixed weakness	Fatigability	Myasthenia gravis against a fixed structural cause
A deficit incongruent with neuroanatomy with positive signs against a deficit that respects it	Internal consistency	Functional neurological disorder against a structural lesion
Cognitive-behavioural decline with myoclonus in a child against chronic disinhibited decline in an adult	Age and the associated sign	Subacute sclerosing panencephalitis against neurosyphilis

MNEMONIC

Course, Leading sign, Unusual features, Exposure and endowment, Systemic and cognitive signs

The recognition spine for any neuropsychiatric stem: **C**ourse and tempo of onset with any fluctuation; the **L**eading neurological sign standing beside the mental-state change; features **U**nusual for the primary disorder being considered; the **E**xposure and inherited endowment the patient brings; and the **S**ystemic and cognitive examination findings a primary formulation would not explain. Read together they spell CLUES, and each is a door back to a treatable disease.

With the pattern recognised, a marks-winning answer follows four moves.

- Name the syndrome and state explicitly that it may be secondary rather than primary.
- Identify the companion neurological sign and use it to localise or to name the likely disease.

- Choose the targeted investigation that answers the specific question, not a reflexive screen.
- Treat on two fronts, the underlying disease and the psychiatric syndrome together.

> In any psychiatric presentation with a neurological companion sign, exclude the treatable disease of the brain before settling on a primary psychiatric diagnosis.

30 · Consolidation and quick review

A chapter that has been read is not yet a chapter that can be retrieved. Reading builds recognition, the comfortable sense that each disease is familiar when its name appears on the page. The examination asks for something harder and less forgiving, which is **free recall**: the ability to generate the right disease, the discriminating feature and the plan from a blank sheet, under time pressure, with the stem giving nothing away. The distance between recognising a syndrome when it is prompted and producing it cold is where marks are won and lost, and it is the distance this closing section is built to close.

Consolidation is the work of turning a read chapter into a small number of durable, retrievable structures. It teaches no new disease. The dedicated sections have already supplied the mechanisms, the investigations and the treatment for each illness in full; the task here is to compress that material into scaffolds that can be reconstructed from memory, and to practise retrieving them so the knowledge survives the weeks between revision and the examination hall. **Consolidation**, in the sense used here, is both the cognitive process by which a memory is stabilised and the deliberate revision method that drives it.

30.1 Why retrieval, not rereading, consolidates this chapter

The single most important decision in revising a dense chapter is how the hours are spent, and memory research and examination experience agree on the answer: memory is built by **retrieval practice**, the act of producing an answer from memory, far more than by rereading the material again. Each time a fact is retrieved with effort, the memory trace that carries it is strengthened and made easier to find next time. Rereading, by contrast, smooths the text into familiarity without building the retrieval route the examination will demand. This asymmetry, that testing a memory improves it more than restudying it, is the **testing effect**, and it is the reason a chapter as long as this one rewards self-testing over passive review.

The practical consequence is to convert reading into questioning at every opportunity. Close the section and state the mechanism, the syndrome and the plan aloud before checking; cover the table and rebuild it; turn each heading into a question and answer it cold. Retrieval that is effortful, and even retrieval that partly fails, consolidates better than fluent rereading, provided the answer is checked and corrected soon afterwards. The discomfort of struggling to recall is not a sign that the method is failing; it is the sensation of the memory being strengthened.

Retrieve	Check	Space	Interleave
GENERATE IT COLD	CORRECT AGAINST THE SOURCE	RETURN AFTER A GAP	MIX THE DISEASES

RULE **You keep what you retrieve, not what you reread.** Plan revision around generating answers from memory and checking them, not around passing your eyes over the text a further time. The chapter is consolidated when each disease can be reconstructed from a blank page, not when it feels familiar on the printed one.

30.2 One retrieval card for every disease: the RECALL template

The chapter contains many diseases, but it asks the same questions of every one of them. That regularity is the reader's advantage, because a single **retrieval template** can be applied to each illness in turn, and the same few prompts will reconstruct any of them. Rather than memorising each disease as an unstructured block, the reader learns one card and fills it repeatedly, which both reduces the memory load and guarantees that no limb of an answer is dropped under pressure.

The template is captured by the word RECALL, chosen because the act it names is exactly the act being practised. Each letter is a prompt to be answered from memory for the disease in hand, and a disease is consolidated when all of them can be produced without the book. The accompanying figure sets the card beside the shared clinical axes as a single revision map.

MNEMONIC

Route, Expression, Clue to organicity, Assessment, Link the treatments, Limit

For any disease in this chapter, retrieve: **R**oute to the brain, the mechanism that carries the disease to the mind, whether direct, immune, infective, systemic and metabolic, or iatrogenic; **E**xpression, the psychiatric and cognitive syndrome the patient shows; **C**lue to organicity, the discriminating red flag that marks the picture as secondary rather than primary; **A**ssessment, the targeted investigation that answers the question the case raises; **L**ink the treatments, the two-front plan that treats the disease and the mind together; and **L**imit, the boundary of what this chapter owns and what it hands to another. Read together they spell RECALL: one card, six prompts, every disease.

The card also disciplines an answer. A candidate who has retrieved all six prompts for the disease in front of them has, without further effort, produced a structured response: the mechanism, the presentation, the reason for suspecting organicity, the investigation, the management and the boundary. The examiner reads that as a formulation rather than a list, which is the difference the marking scheme rewards.

30.3 Retrieve across the mechanisms, not down a list

A list learned in the order it was read is fragile, because recall then depends on starting at the top and running down. Durable retrieval comes from organising the same diseases along more than one axis, so that any entry point reaches the whole set. The first reorganisation is by mechanism. The chapter's diseases reach the mind by a small number of recurring routes, direct involvement of neural tissue, immune and inflammatory injury, infective and post-infective processes,

systemic and metabolic disturbance, and iatrogenic effects of treatment, and retrieving the diseases route by route rather than name by name forces the deeper and more useful connection.

This is **interleaving**: deliberately mixing the diseases during practice instead of studying one to exhaustion before moving to the next. Interleaving feels harder and slower than working through a single disease at a time, and that difficulty is precisely why it works, because it trains the discrimination the examination tests. Retrieving all the diseases that reach the brain by immune injury, or all those that do so by chronic infection, builds the cross-connections that let a reader move sideways from a presentation to its differential rather than only forwards from a name to its facts.

Worked as a drill, this means posing the route as the question. Asked for the immune and inflammatory route, the reader should be able to reach the antibody-mediated encephalitides, the demyelinating disease of multiple sclerosis and the immune hypothesis of psychiatric disorders, each taught in full in its own section. Asked for the chronic infective route, the reader should recover neurosyphilis, the endemic neuroinfections and the subacute and rare chronic infections. The value is not the list itself but the practised movement from a mechanism to the diseases that express it, which is the movement a viva demands.

30.4 Retrieve by the shared clinical axis

The examiner starts from a presentation, so the reader must be able to as well. The second reorganisation runs along the clinical surface, the presenting picture that several diseases share. A stem does not announce a disease; it describes a patient, and the efficient response is to hold the small set of diseases that produce that presentation and then sort them by the single feature that separates them. Retrieving by shared axis rehearses exactly this move. The table below is a retrieval index rather than a teaching table: for each shared presenting axis it names the diseases to be recalled and the prompt that unlocks the discrimination, with the full teaching for each disease living in its own section.

PRESENTING AXIS TO RETRIEVE FROM	DISEASES TO BRING TO MIND	THE PROMPT THAT SORTS THEM
Excess movement with psychiatric change	Huntington disease, Sydenham chorea, the movement lens of Wilson disease, and tics	Course, heredity and age
Slowness and rigidity with a mental-state change	Parkinson disease psychosis and Parkinson disease dementia, and atypical parkinsonism	Drug response, and the timing of cognition and hallucinations
Subacute new psychosis with neurological company	Anti-NMDA receptor encephalitis, the limbic-encephalitis spectrum, and the viral encephalitides	Tempo and the companion neurological sign
Chronic treatable infection presenting psychiatrically	Neurosyphilis, the endemic neuroinfections, and the subacute and rare chronic infections	Exposure, systemic clues and the associated sign
Cognitive, affective or psychotic change on an HIV background	The HIV-associated neurocognitive spectrum, the opportunistic infections, psychiatric comorbidity and drug effects	Which of the four is producing the change
Demyelinating, neuromuscular or functional presentation	Multiple sclerosis, motor neuron disease, myasthenia gravis, and functional neurological disorder	The relapsing course, fatigability or positive signs

Practised in both directions, from an axis to its diseases and from a single disease back to the axis it shares, this index turns the chapter from a sequence of sections into a network that any clue can enter.

30.5 Retrieve the diagnostic boundary

Consolidating this chapter includes consolidating its edges, because a disciplined answer is defined as much by what it hands over as by what it teaches. The **diagnostic boundary** is the set of constructs this chapter names but does not own, and retrieving the boundary cleanly keeps an answer on the disease question actually asked rather than drifting into a neighbouring syndrome and losing time. Each of the following should be recalled as an interface, named where it touches the case and then handed to its owning chapter.

- Pseudobulbar affect, the pathological laughing and crying seen in multiple sclerosis and in motor neuron disease, is named in those diseases but taught in full in the Stroke, TBI, delirium and focal brain syndromes notes.
- Dementia with Lewy bodies, and the temporal rule that separates it from Parkinson disease dementia, belong to the Dementias and neurocognitive disorders notes, while the Parkinson-disease side of that boundary is taught here.
- The diagnostic criteria of Tourette syndrome belong to the Child behavioural, emotional and other disorders notes; the chapter names it as the archetypal tic disorder and teaches the acquired choreas itself.

- Catatonia phenomenology, its structured rating and the lorazepam challenge belong to the Other psychotic disorders, delusional syndromes and catatonia notes, and the chapter teaches only how an antibody-mediated encephalitis is separated from catatonia.
- The primary nosology of depression belongs to the Mood disorders: pharmacotherapy notes; the hepatic and metabolic disease of Wilson disease belongs to the Endocrine and metabolic psychiatry notes; and the standardised cognitive instruments belong to the Psychotherapies, clinical assessment, MSE and nosology notes.

Retrieving these as a short, rehearsed set means the handover sentence arrives ready-made in the examination: name the neighbouring construct, say why it matters to the case, point to its owner, and return to the disease in hand.

30.6 The revision habits that only feel like learning

Several common revision habits produce a strong feeling of progress while building little durable memory, and recognising them is part of consolidating well. The most seductive is the **fluency illusion**: rereading a section until it feels effortless and mistaking that fluency for mastery, when what has actually been trained is recognition of the text, not the ability to generate it from memory. Highlighting, recopying notes and passively watching a topic being explained share the same flaw, feeling productive while leaving the retrieval route unbuilt.

A second habit is massing all the revision of a topic into a single long session close to the examination. Cramming can carry a fact into a test the next morning, but it fades quickly and leaves little for the papers that follow, because memory laid down all at once is denied the spacing that makes it durable. The sense of a productive marathon is real; the retention it buys is not.

PITFALL

Judging readiness by how familiar the material feels rather than by whether it can be produced from a blank page. Familiarity is built by exposure and is a poor predictor of recall; the only honest test of readiness is a cold retrieval attempt with the book closed. A candidate who feels fluent after rereading but cannot reconstruct the RECALL card for a disease has confirmed the illusion, not disproved it, and the remedy is to test earlier and more often, not to reread once more.

TRAP **Confusing recognition on the page with recall in the hall.** Being able to pick the right disease when the material is in front of you is a far weaker skill than generating it from a bare stem, and revision that only ever recognises never trains the harder skill the examination demands. Many candidates discover the gap for the first time in the examination itself, when a syndrome they thought they knew will not come. Practise by producing answers cold, not by reviewing answers already written.

30.7 Space the retrieval and follow the weakness

If retrieval is the engine, **spaced retrieval** is the schedule that makes it last: returning to each topic after a deliberate gap rather than in one continuous block. Spacing works because a memory retrieved just as it begins to fade is strengthened more than one retrieved while still fresh, so allowing a topic to cool a little between attempts is a feature of the method rather than

neglect. Across a revision period, the gaps between visits to a secure topic can lengthen while a shaky one is returned to sooner, so effort follows weakness rather than comfort.

Interleaving and spacing work best together. A practical session mixes several diseases from different parts of the chapter rather than drilling one to exhaustion, and revisits the set after a gap of days rather than minutes. The diseases that fail the retrieval test are flagged and brought forward; the ones that come easily are spaced further out. This concentrates the limited hours of revision on the specific gaps the reader actually has, which no amount of rereading a comfortable section can achieve.

GOOD TO REMEMBER

Build the week around a short daily cold-retrieval test of a mixed handful of diseases, checked and corrected the same afternoon, rather than a long weekly reread. The habit that changes an outcome is closing the book and producing the RECALL card for a handful of diseases from memory, then reviewing only what came out wrong.

30.8 Weight retrieval to the diseases you will meet, and carry the frame in

Not every disease in the chapter carries the same weight at the bedside you will actually work at. Consolidation is finite, so the last decision is where to spend the retrieval that remains, and the honest answer is to weight it toward the diseases that are both common in local practice and dangerous to miss. In much of Indian practice the treatable endemic infections of the nervous system are exactly those diseases, and they reward being retrieved first and most often.

IN INDIA

Where neurocysticercosis, Japanese encephalitis, cerebral malaria, central nervous system tuberculosis, HIV and neurosyphilis remain part of everyday practice, put these treatable causes at the front of the retrieval queue, so that a new psychiatric or cognitive presentation carrying neurological or systemic signs prompts them early rather than late. The clinical decision the weighting protects is to investigate the endemic, potentially treatable disease before settling on a primary psychiatric diagnosis, because several have effective disease-directed treatment while the rest still change supportive care and prognosis, and each is missed when the presentation is read as psychiatric first.

Carried into the examination, the frame collapses to a sequence that survives when detail fails. Meet the stem, decide from its red flags whether the picture is secondary, run the RECALL card for the leading disease, and sort the differential along the shared clinical axis. What consolidation buys is not more facts but faster and more reliable access to the facts already learned, so that under pressure the chapter arrives as a method rather than as a set of half-remembered associations.

Consolidate by retrieval, not rereading: learn one card, RECALL, apply it to every disease, retrieve across mechanism and shared axis, and space the practice until the chapter can be rebuilt from a blank page.

Sources and further reading

The clinical specifics in this chapter are grounded in the standard texts below; where two sources set different thresholds, both are named in the text as a conflict rather than reconciled. Instrument content (the cognitive screens and rating scales) is described and cross-referred, never reproduced.

Sadock BJ, Sadock VA, Ruiz P.
Kaplan and Sadock's Comprehensive Textbook of Psychiatry. 11th ed. Wolters Kluwer, 2024. The primary reference for HIV neuropsychiatry, the movement-disorder psychiatry and the autoimmune and infective encephalitides.

Arciniegas DB, Yudofsky SC, Hales RE (Kaufman).
Kaufman's Clinical Neurology for Psychiatrists. 8th ed. Elsevier. Parkinson disease, Huntington disease, the neuromuscular disorders and the bedside examination.

Brazis PW, Masdeu JC, Biller J.
Localization in Clinical Neurology. 8th ed. Wolters Kluwer. Localisation of the degenerative and demyelinating disorders.

Tasman A, Kay J, Lieberman JA, et al.
Tasman's Psychiatry. 5th ed. Springer. Psychiatric comorbidity in HIV, multiple sclerosis and the systemic diseases.

Geddes JR, Andreasen NC, Goodwin GM (New Oxford).
New Oxford Textbook of Psychiatry. 3rd ed. Oxford University Press. Psychoneuroimmunology and the immune hypothesis.

Johnstone EC, Owens DGC, Lawrie SM, et al.
Companion to Psychiatric Studies. 8th ed. Elsevier. Neurosyphilis, the CNS infections and the hereditary ataxias.

American Psychiatric Association.
Diagnostic and Statistical Manual of Mental Disorders, 5th ed, Text Revision (DSM-5-TR). APA, 2022; and the World Health Organization ICD-11 Clinical Descriptions and Diagnostic Requirements (CDDR).

Taylor D, Barnes TRE, Young AH.
The Maudsley Prescribing Guidelines in Psychiatry. 14th ed. Wiley-Blackwell; and Stahl SM, Prescriber's Guide, Stahl's Essential Psychopharmacology. Cambridge University Press. Prescribing in Parkinson disease psychosis, HIV and neuromuscular disease.

David AS, Fleminger S, Kopelman MD, et al. (Lishman).
Lishman's Organic Psychiatry: A Textbook of Neuropsychiatry. 4th ed. Wiley-Blackwell. The organic and systemic causes of psychiatric presentation across this chapter.

PART VII

VII

Psychopharmacology I: Principles and Core Drug Classes

How psychotropics work and what that predicts: the classes, the receptors, the adverse effects and the reasoning behind a first choice.

01 Receptor pharmacology and pharmacokinetic principles

02 Antipsychotics

03 Antidepressants

04 Mood stabilisers

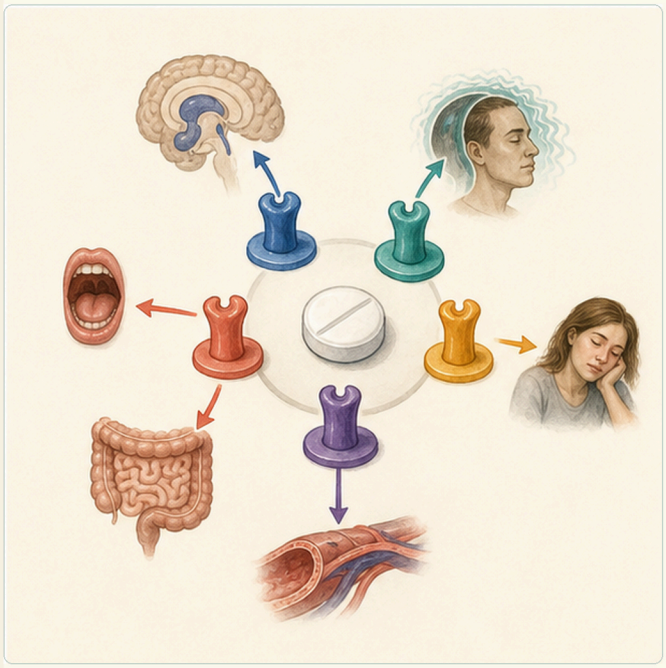
05 Anxiolytics, hypnotics and other agents

06 Consolidate and apply

07 Sources and further reading

VISUAL 7.1

A receptor profile predicts a burden profile



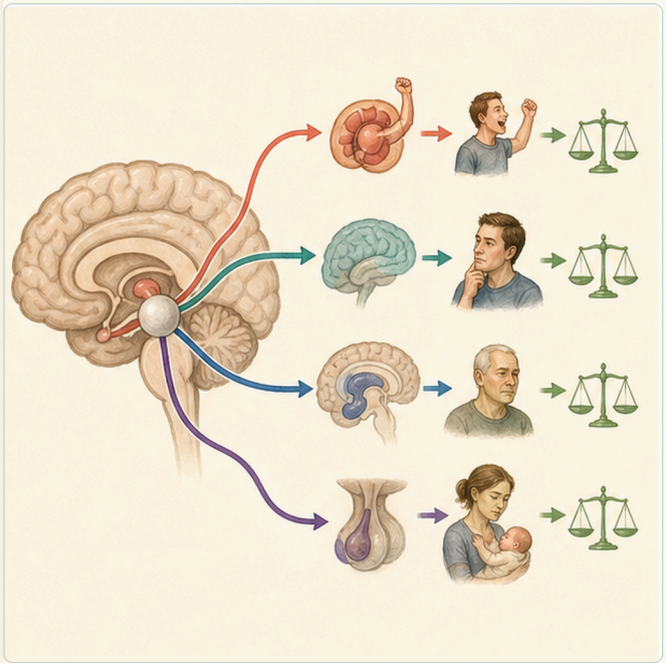
- D2
- 5-HT2A
- H1
- M1
- Alpha-1

A psychotropic receptor profile helps predict therapeutic and adverse-effect patterns, but exposure and individual context still matter.

Source lock: day-36.final.html#sec-t02

VISUAL 7.2

One dopamine mechanism, four pathway consequences



- Mesolimbic
- Mesocortical
- Nigrostriatal
- Tuberoinfundibular
- Occupancy
- Benefit-harm balance

Dopamine-receptor effects differ by pathway, producing distinct therapeutic and adverse consequences from one pharmacological action.

Source lock: day-36.final.html#sec-t08

VISUAL 7.3

Lithium safety is a monitoring system



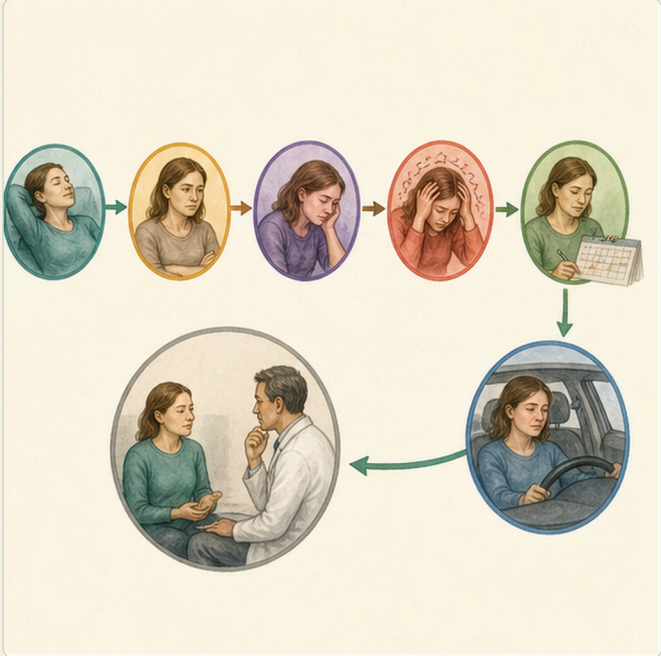
- Benefit
- Renal and thyroid
- Timed level
- Toxicity
- Interactions
- Hydration and illness

Lithium's clinical value depends on dose, timed levels, organ monitoring, interaction control and early recognition of toxicity.

Source lock: day-36.final.html#sec-t19

VISUAL 7.4

Benzodiazepine exposure becomes a withdrawal problem



- Short-term relief
- Tolerance
- Dependence
- Withdrawal
- Planned taper
- Residual impairment

Benzodiazepine benefit, tolerance, dependence, impairment and withdrawal are exposure- and duration-dependent parts of one prescribing problem.

Source lock: day-36.final.html#sec-t23

PAPER IV

Psychopharmacology I: principles and core drug classes

Five bands . one complete notes document

Receptor pharmacology and pharmacokinetic principles

1 · Psychopharmacology as a system: what this chapter owns, and what it points to

A psychotropic is prescribed for a disorder, but it acts on receptors, transporters and enzymes that have no knowledge of the diagnosis. The benefit and the adverse effects arise from the same molecular actions, distributed across different tissues and pathways. This is why one molecule can be reached for in a depressive episode, an anxiety disorder, a pain syndrome and an eating disorder: the indication changed, the pharmacology did not. A clinician who has learned drugs only through the disorders they treat can say what to give but not why it works, cannot anticipate the next adverse effect, and cannot choose between two agents said to be interchangeable.

These notes therefore study the drug class as an object in its own right. The questions are pharmacological throughout: what does this class bind, what does binding there produce, how does the body handle the molecule, how much room lies between the exposure that helps and the exposure that harms, and what separates the members of a class from one another. The disorders are not absent, but they are the setting rather than the subject, and the chapters that own each clinical decision are named here so that their material is not repeated. Marks in this territory are given for mechanism that predicts, not for lists that recall.

1.1 The lens: pharmacology as the subject, not the setting

The organising lens of this chapter is **the systematic pharmacology** of the psychotropic classes, and it carries three commitments. The first is that the unit of study is the **drug class**, defined by a shared mechanism rather than by a shared indication, so that what is learned about one member transfers to the rest and the exceptions become informative instead of arbitrary. The second is that the adverse-effect profile of an agent is a prediction derived from its pharmacology, not a list to be memorised drug by drug. The third is that choosing between members of a class is a pharmacological comparison: agents that share an indication are separated by what they bind, by how the body handles them and by their margin of safety, and those are the axes on which a rational choice is made.

What the lens deliberately excludes is the treatment of any disorder. Which antipsychotic to start in a first episode, when to switch an antidepressant, how long to continue prophylaxis and what to do when a patient does not respond are clinical questions answered elsewhere: in the Schizophrenia: management, antipsychotics and adverse effects notes, in the Mood disorders: pharmacotherapy notes, in the Dissociative, conversion and anxiolytic/hypnotic pharmacology

notes and in the ADHD and specific learning/communication disorders notes. Nothing here replaces those decisions; it supplies the reasoning on which they rest.

TRAP **Answering a pharmacology question with a treatment algorithm.** Asked why a class produces the adverse effects it does, or how two members of a class differ, candidates reproduce the management of the disorder the class happens to treat. The algorithm may be entirely correct and still earn nothing, because the question was about the drug. Name the receptor action, the kinetic property or the safety margin the difference rests on, and only then say what it means at the bedside.

1.2 Four questions that characterise any psychotropic

Every agent in this chapter can be characterised by the same four questions, asked in the same order; the sections that follow answer them for each class in turn. The order matters, because each question constrains the next.

The first asks what the drug binds. A **receptor binding profile** is the set of receptors, transporters and enzymes the molecule engages, together with the relative strength of each engagement, and it is the closest thing pharmacology has to a fingerprint. The second asks what the drug does where it binds, because occupying a target is not a single behaviour: a molecule may reproduce the natural signal, block it, produce a submaximal signal that acts as a brake where transmission is high and as a support where it is low, or suppress activity below the unstimulated baseline. That is the **mode of action**, and it changes what an identical binding profile means clinically. The third asks what the body does to the drug: how it is absorbed, how it is cleared, and therefore how long its effect persists after a dose and after the last dose. The fourth asks how much room separates the exposure that helps from the exposure that harms.

Profile	Action	Fate	Margin
WHAT IT BINDS	WHAT IT DOES THERE	WHEN IT ACTS AND FOR HOW LONG	HOW CLOSELY TO WATCH

The four answers together predict most of what a prescriber needs: the likely benefit, the likely adverse effects, when to expect either, what happens on stopping, and which member of a class suits a particular patient. The biochemistry underneath the first two questions, meaning the transmitter systems themselves, their synthesis and reuptake, their pathway anatomy and their signalling cascades, belongs to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes. Here the receptors are treated as targets whose engagement has consequences a patient can feel and a clinician can foresee.

1.3 The receptor profile predicts the adverse-effect signature

The most portable idea in this chapter is that the **adverse-effect signature** of an agent is readable from its binding profile before any patient has been exposed to it. A drug is prescribed for what it does at one target, and it also acts at every other target it binds; those other actions are what the patient reports as side effects. Sedation, dry mouth, constipation, blurred vision, postural dizziness, weight gain, movement disturbance and endocrine change are not a random accompaniment of psychiatric treatment. Each is the predictable consequence of engaging a

particular receptor in a particular tissue, and the same receptor action produces the same complaint whichever class the molecule belongs to.

Two consequences follow, and both are worth carrying into an examination. The first is that an unfamiliar drug is not an unknown drug: given its binding profile, its probable burden can be stated in advance, which is precisely what an examiner who names a recent agent is testing. The second is that **off-target action** is a design axis rather than a defect to be regretted. Much of what separates an older from a newer member of the same class is that the newer one binds fewer targets, and much of the residual usefulness of the older one is that some of those extra actions are therapeutic in the right patient. The section on reading a receptor profile builds that map in full; what matters at this stage is the habit of asking what else the molecule binds.

RULE **Every predictable adverse effect of a psychotropic is a pharmacological action somewhere other than where the benefit is.** Ask what else the molecule binds, in which tissue, and what that tissue does when the receptor is engaged. An adverse effect that can be derived is an adverse effect that can be anticipated, warned about in advance, and avoided altogether where a member of the same class lacks that action.

1.4 Kinetics decides timing, and timing decides most errors

Binding explains what a drug does; kinetics explains when. The property carrying most of that work is the **elimination half-life**, the time taken for the concentration to fall by half once input stops. It governs three things a prescriber uses constantly: how long a drug takes to reach a stable concentration on a fixed dose, when a measured concentration means anything, and how abruptly the drug leaves the patient when it is stopped.

A drug given repeatedly accumulates until the amount eliminated between doses equals the amount given, and **steady state** is that plateau. Two rules of practice follow from it. A concentration measured before the plateau describes a rising curve rather than the exposure the patient will actually run on, so acting on such a value invites a dose change that the settled level would never have justified. And the speed of offset on stopping is a kinetic property rather than a moral one: an agent cleared quickly falls away fast enough to generate discontinuation phenomena, while an agent cleared slowly tapers itself. Whether a drug needs a planned withdrawal is therefore predictable from its kinetics before any patient has complained of anything.

PITFALL

Checking a concentration too early and then changing the dose on the strength of it. A level drawn before the drug has plateaued, or drawn at an inconsistent point in the dosing interval, reports a value that is real and uninterpretable. The dose is adjusted against an artefact, the next level is drawn before the new dose has settled either, and the patient ends up titrated on noise. Know when the drug plateaus, sample consistently in relation to the dose, and treat a single unexpected value as a reason to repeat the test rather than to act on it.

1.5 The therapeutic index decides how closely to watch

The **therapeutic index** expresses the distance between the exposure that produces the wanted effect and the exposure that produces harm. Where that distance is wide, ordinary variation in absorption, hydration, renal function or a co-prescribed drug moves the patient about within a tolerant range, and clinical observation alone is sufficient supervision. Where it is narrow, the same ordinary variation carries the patient out of treatment and into toxicity, and observation is not enough, because the early features of toxicity in a narrow-index drug can be mistaken for the illness being treated or for the drug's own familiar adverse effects.

This is the pharmacological reason a small number of psychotropics are monitored by measurement while most are not. Monitoring is not a ritual attached to old drugs; it is what a narrow index costs. The same reasoning explains why dehydration, a change in renal function, an added diuretic or an interacting agent matter disproportionately for those drugs, and why teaching the patient to recognise those precipitants is part of the prescription rather than an optional courtesy. The target concentrations themselves, and the wider framework of therapeutic drug monitoring in psychiatry, belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes. What is taught here is why a drug earns that scrutiny in the first place.

GOOD TO REMEMBER

Before writing any psychotropic prescription, be able to say four things out loud: what the drug binds besides its intended target, whether it is cleared quickly or slowly, whether its margin of safety is wide or narrow, and what could plausibly change in this particular patient to shift the exposure. A prescriber who can answer those has already predicted most of what will happen next, including the complaint the patient will bring to the following appointment.

1.6 Read the chapter as a map of drug territories

The remaining sections are arranged by pharmacological family, and the arrangement is deliberate. The principles come first because they are the tools used on everything after them, and each territory is then examined in the same way, so the reader acquires one method applied repeatedly instead of several unrelated bodies of fact. The table states what each territory is for; the sections supply the depth.

TERRITORY	THE PHARMACOLOGICAL QUESTION IT ANSWERS	WHAT THE SECTIONS SUPPLY
Receptor pharmacology and kinetic principles	How is any psychotropic characterised at all?	Reading a binding profile, what agonism, antagonism and partial agonism change for a prescriber, half-life and steady state, the cytochrome substrate map in practice, and which agents need a measured level
Antipsychotics	What does a single occupancy mechanism produce across several pathways?	What the typical and atypical labels do and do not predict, occupancy and its consequences, the movement-disorder spectrum, the clozapine outlier, injectable release kinetics, and the within-class comparison
Antidepressants	What separates families that share an indication?	One mechanism axis across the families, the differences inside the selective reuptake inhibitors, the distinct receptor actions of the newer agents, tricyclic receptor promiscuity, and the interaction burden of enzyme inhibition
Mood stabilisers	What does a narrow margin of safety demand of a prescriber?	Lithium as the archetypal narrow-index drug, the kinetic precipitants of toxicity, valproate as a class member, enzyme induction and autoinduction, and a titration that pharmacology forces
Anxiolytics, hypnotics and other agents	Why do tolerance, dependence and withdrawal follow from mechanism?	Benzodiazepine subunit pharmacology, the hypnotics compared, the anxiolytics that behave unlike a benzodiazepine, the cognitive enhancers named as a class, and stimulant onset, duration and diversion

Read the map in both directions. Given a drug, it says which pharmacological questions to ask of that drug. Given a question about mechanism, adverse effect or choice, it says where the answer is taught. The value of the arrangement is that the method stays constant while the classes change, so a reader who has worked through one territory already knows how to interrogate the next.

IN INDIA

Within-class choice is the decision most worth rehearsing, because the question in an ordinary outpatient clinic here is usually which member of a class to use rather than whether to use the class at all. Two habits follow from the pharmacology. Make the comparison on binding, kinetics and margin of safety rather than on prescribing habit or familiarity. And where a drug's clearance depends on an enzyme whose activity varies with ancestry, start at the lower end of the class range and titrate against response and adverse effect rather than to a fixed schedule. The sections on the cytochrome substrate map and on within-class comparison carry the detail that makes each of those decisions specific.

1.7 Know the boundary and hand over cleanly

Competence includes knowing where a question stops belonging to pharmacology. Naming the **handover boundary** and pointing to the chapter that owns the material is shorter, more accurate and better rewarded than reproducing a neighbouring chapter's algorithm from memory, and it keeps an answer on the question actually asked. The handovers that matter most are these.

- Choosing, switching and continuing antipsychotic treatment in a patient with schizophrenia, and the clinical management of its adverse effects, belong to the Schizophrenia: management, antipsychotics and adverse effects notes.
- Selecting, sequencing and continuing an antidepressant or a mood stabiliser for a depressive or a bipolar episode belongs to the Mood disorders: pharmacotherapy notes.
- The clinical use of the anxiolytics and hypnotics in anxiety and insomnia belongs to the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes, while benzodiazepine dependence as a diagnosable disorder, with its assessment and its detoxification, belongs to the Substance use disorders notes.
- Indication-linked stimulant prescribing and titration in attention deficit hyperactivity disorder belongs to the ADHD and specific learning/communication disorders notes.
- Valproate, carbamazepine and lamotrigine used as antiepileptic drugs, with the titration and precautions specific to epilepsy, belong to the Epilepsy and psychiatry notes.
- Interactions that cross drug classes, prescribing in physical illness and across the lifespan, monitoring that spans classes, therapeutic drug monitoring target ranges and the psychotropic emergencies all belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

A handover sentence does real work when it names the neighbouring construct, says why it matters to the patient in front of you, points to where it is taught, and then returns to the pharmacological question. Notice the principle organising the last of those handovers in particular: this chapter compares agents within a class, and the moment a comparison crosses class boundaries it becomes a different chapter's subject. That single distinction resolves most uncertainty about where a given piece of prescribing knowledge belongs.

1.8 Carry the PROFILE scaffold into essays and vivas

Under examination pressure the method has to survive when the detail deserts you. The spine holding this chapter together is **PROFILE**. It is a teaching scaffold rather than a validated instrument, and it will structure an answer about any class in these notes, whether the question is a mechanism, an adverse effect, a comparison between two agents or an unfamiliar drug the examiner has just named.

MNEMONIC

PROFILE

Place the drug in its class and say what the class shares. **R**eceptors: what does the molecule bind, and what else does it bind. **O**ccupancy behaviour: does it reproduce, block, partially support or suppress the signal, and what does that change. **F**ate in the body: absorption, clearance, half-life and steady state. **I**ndex: how wide is the margin between benefit and harm, and does the drug therefore need a measured level. **L**iability ranked within the class: sedation, weight, endocrine and movement burden compared across members. **E**lsewhere: name what belongs to another chapter and hand it over.

For a short answer, compress PROFILE into class, mechanism, principal adverse effects and the one comparison that decides the choice. For a long essay, expand each letter and draw the specifics from the owning section. In a viva, make the derivation audible: state what the drug binds, derive the effect from it, then say what you would tell the patient to expect and when. The examiner is listening for a clinician who reasons forward from mechanism, because that clinician stays safe with an agent no textbook has yet described.

PROFILE: place the drug in its class, read the receptors it binds, ask what it does when it binds them, follow its fate in the body, weigh the therapeutic index, rank liability within the class, then hand over what belongs elsewhere.

2 · Reading a receptor profile: predicting the adverse-effect signature of any agent

An adverse-effect list is not something to be memorised agent by agent, it is something to be derived. Almost every psychotropic in routine use binds more than one receptor family, and the receptors it engages beyond its intended target are where the unwanted effects are generated. Learn which receptor action produces which complaint, and the tolerability profile of any molecule can be reconstructed from its binding data, including molecules that did not exist when the reader trained. That is the whole content of this topic, and it is the highest-yield reasoning step in psychopharmacology, because it replaces an unbounded memorisation task with one short map applied over and over.

Examinations exploit the same structure directly. A question that names an unfamiliar agent, gives its receptor affinities and asks which adverse effect to counsel the patient about is testing whether the candidate can run the map forwards; a question that gives a complaint and asks which receptor action explains it runs the same map backwards. Both are answerable from one

table. What follows walks it receptor family by receptor family: what each blockade predicts, which entries on the list are protective rather than costly, why the same receptor produces unrelated complaints in different tissues, and why the class of receptor a drug acts on predicts how fast its clinical effect arrives. The biochemistry underneath, meaning transmitter synthesis, reuptake, pathway anatomy, receptor subtype classification and the second-messenger cascades, belongs to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes and is assumed here.

H1	M	alpha-1	D2
SEDATION, AND PART OF THE WEIGHT GAIN	DRY MOUTH, VISION, BLADDER, BOWEL	DIZZINESS, HYPOTENSION, SYNCOPE	MOVEMENT DISORDER, ACUTE AND TARDIVE

2.1 The profile, not the label, predicts the patient

A drug's class name records the indication it was developed for. Its **binding profile** records what it will do to the person who swallows it, and the two are different documents. One molecule ordinarily engages several receptor families at once, which is why a signature of unwanted effects has to be read off the profile as a whole rather than inferred from the name on the box. An agent filed under a class heading may sit at the promiscuous end of that class or the selective end, and nothing in the heading tells the reader which.

Two consequences follow at once. The first is that agents sharing a therapeutic class can differ sharply in tolerability, because a shared primary target says nothing about what else each molecule touches; the within-class comparison of the antipsychotics has its own section in these notes. The second is that an adverse effect is not evidence of a drug misbehaving. It is evidence of the drug doing precisely what its affinity data predicted, at a receptor nobody wanted engaged. That reframing is not cosmetic: it converts the conversation from apology after the event into warning before it, and a warned effect is tolerated far better than a surprising one.

RULE **The receptor profile generates the adverse-effect signature; the class label does not.** Read affinities, not headings. This is why a newly licensed molecule can be counselled about competently on the day it arrives, why two agents with the same primary target can behave differently at the bedside, and why an adverse effect appearing in a patient is a finding to be located on the profile rather than a mystery to be investigated.

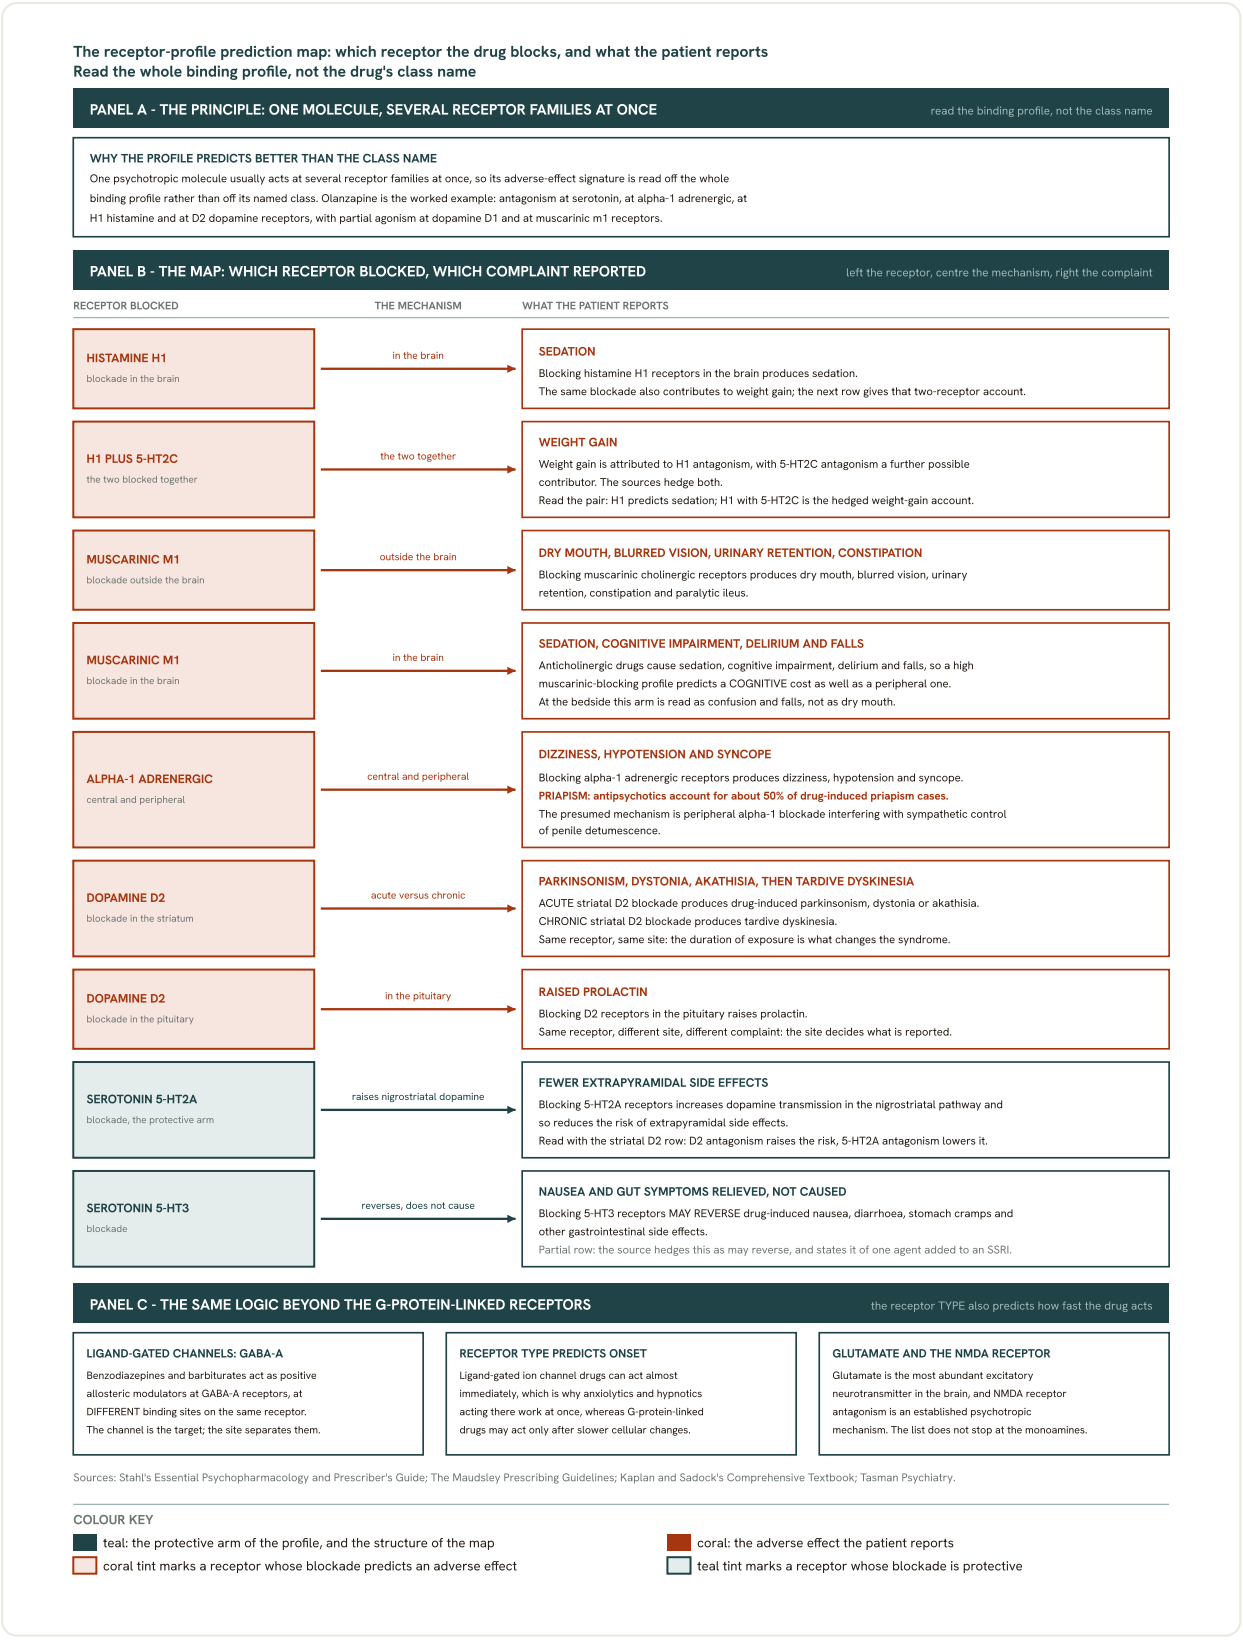


Fig 36.1 **The receptor-profile prediction map.** The chapter's organising figure: the adverse effect a patient reports is read off the drug's whole binding profile, not off its class name. H1 blockade predicts sedation, H1 with 5-HT2C weight gain, muscarinic blockade a peripheral and a cognitive cost, alpha-1 blockade dizziness, hypotension and syncope, striatal D2 blockade the movement disorders and pituitary D2 blockade a prolactin rise, while 5-HT2A and 5-HT3 blockade sit on the protective side of the same profile. (Stahl's Essential Psychopharmacology and Prescriber's Guide, The Maudsley Prescribing Guidelines, Kaplan and Sadock, and Tasman Psychiatry.)

2.2 Histamine: sedation, and part of the weight problem

The most immediately visible entry on the map is histaminergic, because the patient reports it within days and the family notices it before the patient does. Central **histaminergic H1 blockade** is what makes a patient sedated, and the same receptor action is also implicated in weight gain. Sedation of this origin is not a marker of a drug being powerful or of the illness responding; it is a marker of affinity at a receptor with no relationship to the therapeutic target, and it appears early because it requires no adaptive change to develop.

Weight gain is the more consequential part of the histaminergic contribution, and it is the part the sources hedge. Histamine blockade is recorded as causing sedation and possibly weight gain, and the fuller account is a joint one: H1 antagonism together with antagonism at **5-HT2C** may explain some aspects of drug-induced weight gain. Read 5-HT2C in a profile, then, as a further contributor alongside the histaminergic effect rather than an independent predictor on its own. Reproduce the hedge when you answer, because the mechanism of drug-induced weight gain is not settled, and a candidate who states any one receptor flatly as its cause has claimed more than the sources do.

TRAP **Assuming the map is one-to-one, with each effect owned by exactly one receptor.** It is not one-to-one in either direction. Weight gain is attributed to more than one receptor action, each stated as a possible contributor rather than a settled single cause, and sedation has more than one route on to the map, since histaminergic blockade and central muscarinic blockade both produce it. Candidates lose marks by naming one receptor flatly as the cause of an effect the sources attribute jointly and hedge, and clinicians lose ground by withdrawing the wrong agent when a sedated patient is taking two drugs that reach the same complaint by different receptors.

GOOD TO REMEMBER

A profile with strong affinity at H1 predicts a patient who will be sleepy in the first days and heavier at review. Both belong in the counselling given at the point of prescribing, not in the explanation offered after the patient has stopped the drug. Weighing the patient before the first dose costs nothing and makes the second conversation possible.

2.3 Muscarinic blockade: the visible bill and the invisible one

Muscarinic cholinergic antagonism is the widest-ranging entry on the list, because the receptor is distributed through the eye, the salivary glands, the bladder and the gut as well as the brain. Peripherally, **antimuscarinic blockade** predicts a dry mouth, blurred vision, urinary retention and constipation, with paralytic ileus sitting at the far end of the same continuum rather than constituting a separate phenomenon. The gradient matters clinically: constipation on a strongly antimuscarinic agent is not a minor inconvenience to be managed later, it is the early portion of a curve whose distal end is an obstructed bowel.

The central portion of the bill is the one routinely underweighted, and it is the portion that damages function. A high muscarinic-blocking profile carries an **anticholinergic burden** whose consequences are sedation, cognitive impairment, delirium and falls, so the profile predicts a cognitive cost as well as a peripheral one. Those effects are more severe in older patients with

dementia, and prescribing in the elderly belongs to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes. Those notes also own anticholinergic toxicity as an acute emergency; what belongs here is only the prediction running from receptor to symptom.

> **PITFALL**
>
> Estimating anticholinergic burden from the dry mouth. The peripheral symptoms are the visible marker of the load, not its measure, and a patient who volunteers none of them may still be carrying a cognitive cost that neither the patient nor the family will attribute to the tablets. Ask specifically about concentration, recent memory and unsteadiness in anyone on a strongly antimuscarinic profile, and ask the informant rather than only the patient.

2.4 Alpha-adrenergic blockade: the postural signature

The adrenergic entry is the one whose consequences are mechanical rather than experiential, and it is the one most likely to injure a patient outright. **alpha-1 adrenergic antagonism** predicts dizziness, hypotension and syncope, a sequence that runs from a symptom the patient tolerates to an event that puts them on the floor. Because the mechanism is loss of reflex vasoconstriction rather than a central effect, the symptom is provoked by posture and by anything that reduces circulating volume, and it is worst at the start of treatment and after any increase in dose, before the circulation has adapted.

The same receptor action generates one effect the patient will not raise unprompted. **Priapism** is a recognised consequence of adrenergic blockade, and the attribution is not marginal: **about 50 per cent** of drug-induced cases are laid at the door of antipsychotics. The presumed pathophysiology is peripheral alpha-1 adrenergic blockade interfering with the sympathetic control of penile detumescence, which is the same molecular action responsible for the postural symptoms, expressed in a different vascular bed. It is the clearest instance on this map of an adverse effect that must be asked about, because it will not be volunteered. Prescribing in the presence of physical illness belongs to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

2.5 Dopamine: one receptor, different tissues, different timescales

Nothing on this map demonstrates the importance of tissue as clearly as the dopamine entry. Blockade at D2 receptors produces entirely unrelated clinical pictures depending on where in the body the blockade occurs, and the receptor name alone therefore predicts nothing useful. In the striatum, the effect is further split by time. Acute striatal D2 blockade produces drug-induced parkinsonism, dystonia or akathisia, whereas chronic blockade of the same receptors at the same site produces **tardive dyskinesia**. One receptor, one tissue, opposite ends of the exposure curve, different syndromes. The spectrum of these movement disorders and how each is recognised has its own section in these notes; their clinical management in a treated patient belongs to the Schizophrenia: management, antipsychotics and adverse effects notes.

Move the same blockade to the pituitary and the output changes completely: blocking D2 receptors there raises **prolactin**. Dopaminergic tone at that site is inhibitory, so antagonism removes a brake rather than adding a stimulus, and the endocrine consequence follows without

any involvement of the motor system. The pathway-by-pathway account of dopamine occupancy is taught in its own section of these notes and is not repeated here. What this topic takes from it is the general lesson: a receptor plus a tissue is a prediction, whereas a receptor alone is only part of one.

2.6 The serotonergic entries: where the accounting changes sign

Everything to this point has been a cost. The serotonergic entries are where the map records a benefit, and this asymmetry is worth learning explicitly, because candidates who have absorbed the idea that receptor blockade means adverse effect answer these questions backwards. Blockade at **5-HT2A** increases dopaminergic transmission in the nigrostriatal pathway, and since extrapyramidal effects arise from insufficient dopamine at that site, the same molecular action lowers the risk of extrapyramidal side effects. Dopamine antagonism raises that risk; added 5-HT2A antagonism lowers it. A profile combining the two therefore behaves differently at the motor level from dopamine antagonism alone, which is the pharmacological content of a distinction that classification labels only gesture at. What the typical and atypical labels do and do not predict has its own section in these notes.

The other protective entry is at **5-HT3**. Blockade there may reverse drug-induced nausea, diarrhoea, stomach cramps and other gastrointestinal effects, so a molecule carrying that action can be expected to sit more comfortably in the gut than its mechanism at other receptors would suggest. The evidence should be held at the strength the source states it: the observation comes from adding an agent with 5-HT3 antagonism to a serotonin reuptake inhibitor, and the source hedges the effect rather than asserting it. The receptor-level generalisation is a reasonable reading, not a quoted certainty.

That receptor is also, on this list, the one whose blockade the standard sources record only as protective. A symmetrical adverse effect attributable to it was searched for and not found, and none is supplied here. An honest blank on a map is worth more than a plausible invention, because the invention is what a candidate reproduces under pressure.

2.7 Beyond the G-protein-linked receptors: channels, and speed

The receptors walked through so far signal through G proteins, but not every psychotropic target does. The benzodiazepines and the barbiturates are the worked example: each group acts as a **positive allosteric modulator** at GABA-A receptors, and each does so at a different binding site on that same receptor. Different sites on a shared receptor is precisely why drug groups with a similar gross clinical effect are not pharmacologically interchangeable. Benzodiazepine class pharmacology has its own section in these notes; dependence and withdrawal as a disorder belong to the Substance use disorders notes, and their use as clinical treatments for anxiety and insomnia to the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes.

The class of receptor also predicts how quickly the patient notices anything, which is the part of this topic patients ask about most often. Because an **ionotropic receptor** changes ion flow the instant it is bound, a drug acting there can have an almost immediate effect, which is why anxiolytics and hypnotics working at these receptors are felt at once. A drug acting at a G-

protein-linked receptor may instead produce its clinical effect only after slower changes in cellular function have been set in motion through the signal transduction cascade, so a delay between the first dose and the benefit is a property of the target rather than a failure of the drug or of the patient. Onset speed, in other words, is one more thing readable off the profile before treatment starts.

Glutamate completes the target list. It is the most abundant excitatory neurotransmitter in the brain, and **NMDA receptor antagonism** is an established psychotropic mechanism rather than a laboratory curiosity. The agents built on it belong to the antidepressant sections of these notes; what belongs here is only the recognition that the prediction map is not confined to the monoamine and cholinergic receptors that dominate older profiles.

2.8 Running the map on a molecule you have never seen

The test of a map is whether it works on unfamiliar ground. Given a new agent and a table of its receptor affinities, a short sequence of questions extracts most of the counselling list before any clinical experience with the drug exists.

- Which receptors does it block, and how strongly? Affinity ranks the predictions, and the strongest action away from the intended target is usually the complaint the patient brings first.
- Where does each of those receptors sit? The same blockade in brain, gut, bladder and pituitary produces unrelated pictures, so the tissue is part of every prediction.
- Does any entry on the profile offset another? A protective action changes the net expectation and is often why agents sharing a primary target diverge in the clinic.
- Is the target a channel or a G-protein-linked receptor? That decides whether the patient feels something after the first dose or after considerable time on treatment.

The accompanying figure sets out the prediction map itself. The table below works the same relationships in the direction the clinic actually poses them, from the complaint the patient reports back to the receptor action that produced it, which is the harder direction and the one examinations increasingly prefer.

WHAT THE PATIENT REPORTS	RECEPTOR ACTION TO SUSPECT	WHAT SEPARATES IT
Sedation	Histaminergic H1 blockade, or central muscarinic blockade	More than one route: look for antimuscarinic features before attributing it to histamine
Weight gain	H1 together with 5-HT2C blockade	The joint account, hedged at source; no single receptor is the settled cause
Dry mouth, blurred vision, urinary retention, constipation	Peripheral muscarinic blockade	Paralytic ileus sits at the far end of the same continuum
Cognitive slowing, delirium, falls	Central muscarinic blockade	Ask directly; the peripheral symptoms may be absent

WHAT THE PATIENT REPORTS	RECEPTOR ACTION TO SUSPECT	WHAT SEPARATES IT
Postural dizziness, syncope	alpha-1 adrenergic blockade	Hypotension is its measurable form; check lying and standing
Priapism	alpha-1 adrenergic blockade, peripherally	Ask about it; the patient will not raise it
Parkinsonism, dystonia, akathisia	Acute striatal D2 blockade	Appears with recent exposure or a recent increase
Involuntary movements emerging late	Chronic striatal D2 blockade	Tardive dyskinesia; the exposure history separates it
Raised prolactin	Pituitary D2 blockade	Tissue, not receptor, distinguishes it from the motor effects
Nausea, diarrhoea and cramps easing	5-HT3 blockade	A protective entry, and no adverse counterpart is recorded
Less extrapyramidal effect than dopamine blockade alone predicts	5-HT2A blockade	The offsetting entry on the profile

MNEMONIC

HAM, then D, then the serotonin credit

Histamine, **A**drenergic alpha blockade and **M**uscarinic blockade account for most of the routine tolerability complaints of a broad-profile agent: sedation with part of the weight gain, dizziness with hypotension and syncope, and the dry mouth, blurred vision, retention and constipation with their central cognitive cost. Add **D**opamine for the movement and prolactin effects. Then remember that the serotonergic entries can sit on the credit side of the ledger rather than the debit side.

Finally, the map states its own limits, and a candidate who states them is a candidate who understands it. A receptor profile predicts effects that are receptor-mediated and broadly related to exposure. It does not predict idiosyncratic reactions, immunologically mediated events, organ toxicity or the rare catastrophic outcomes, none of which have an affinity column to be read off, and it says nothing about how the drug is handled by the body before it reaches any receptor at all. Reading the profile is the first step of the reasoning about any agent, and a reliable one; it is not the whole of it.

Blockade of histaminergic, muscarinic, alpha-adrenergic and dopaminergic receptors accounts for most of what patients complain about; the serotonergic entries are the arm of the profile that can protect.

3 · Agonism, antagonism and partial agonism as a prescribing variable

Whether a receptor drug helps, does nothing, or moves transmission the wrong way can turn entirely on how much of the natural transmitter is already sitting at that receptor. Agonist,

antagonist and partial agonist are usually learnt as vocabulary and then used as a switch: one turns the receptor on, one turns it off. That reading does not survive contact with a prescription. The same molecule at the same dose can raise transmission in one patient and lower it in another; a drug classified and marketed as an antagonist may be doing something an antagonist cannot do; and a partial agonist can be pushed to the top of its licensed range without the patient noticing any further effect at all. Each of those is a predictable consequence of where a molecule sits on one continuum, and each is a question an examiner can put in a single line. What follows treats that continuum as a prescribing variable: what each position does to signal transduction, what the position predicts about a given patient's response, and where the prediction breaks down. The receptor science underneath it, meaning occupancy mathematics, the receptor-regulation cascade and the molecular basis of up-regulation, down-regulation, tolerance and sensitisation, belongs to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes and is assumed here.

4	Maximal	Ceiling	Below
NAMED POSITIONS ON ONE CONTINUUM	FULL AGONIST OUTPUT AT THE RECEPTOR	PARTIAL AGONIST OUTPUT, FIXED BY THE MOLECULE	INVERSE AGONIST, UNDER THE UNOCCUPIED BASELINE

3.1 One continuum, four positions

Ligands acting at a G-protein-linked receptor do not sort into two boxes. They occupy **four** named positions along a single continuum, ordered by how much signal transduction each produces: **the agonist spectrum** begins at the full agonist, passes through the partial agonist and then the silent antagonist, and ends at the inverse agonist. Every clinical prediction that follows is a consequence of that ordering, so what is worth holding is the sequence itself rather than four separate definitions memorised apart from each other.

At the top sits the **full agonist**, which produces the conformational change that turns on second-messenger synthesis to the greatest extent the receptor permits. It is usually the naturally occurring neurotransmitter itself, although some drugs act as completely as the transmitter does. That qualification is easy to skim past and expensive to skim past. The reference point at the top of the spectrum is not an abstraction on a graph; it is a molecule already present in the synapse of the patient in front of you, at a concentration that varies with brain region, with physiological state and with illness. Because every lower position is defined against it, every lower position is implicitly defined against a quantity that cannot be measured in clinic and has to be inferred from the clinical picture.

Two habits of thought follow, and they are the reason this material is a prescribing topic rather than a definitions topic. The first is to stop asking what a drug does and start asking what it does in this system, in this patient, in this phase of illness. The second is to treat the four labels as points on a scale of output rather than as separate species of drug: they differ in degree along one axis, and the boundaries between them are where the interesting clinical behaviour lives.

3.2 The silent antagonist has no pharmacology of its own

A **silent antagonist** produces a conformational change that causes no change in signal transduction, and leaves any resting output the receptor already had exactly where it was. It is neutral in the strict sense: it occupies the receptor and does nothing there, which is precisely why such compounds are also described as silent. The definition takes one line and its consequences take the rest of this section.

The first consequence is that an antagonist's clinical effect is entirely borrowed. Given into a system where no agonist is present, it changes nothing whatever. Given into a system where a full agonist is present, it blocks that agonist and still contributes nothing itself. What the prescriber reads as the drug's potency is therefore a statement about the patient's transmitter tone at least as much as about the drug, which is why an identical exposure to the same blocker looks dramatic in one clinical state and inert in another. Nothing about the tablet has changed between those two patients.

The second consequence is that an antagonist reverses the whole spectrum rather than agonism alone. It blocks the actions of everything on the spectrum and simply returns the receptor to the conformation that exists when no agonist is present, and from that one fact the full list of what it undoes can be derived without memorising it:

- a full agonist, whose maximal signalling is removed and whose effect the blocker was designed to oppose;
- a partial agonist, whose actions at that receptor an antagonist also blocks, so the two compete for the same site and what a pairing of them produces depends on which is occupying it;
- an inverse agonist, whose suppression below the resting baseline is also lifted, which is the least intuitive of the three and the one most often missed.

The third consequence is the one candidates reverse under time pressure. An antagonist is not the opposite of an agonist. It is the absence of an agonist, reimposed on the receptor by physically occupying it. The genuine opposite of agonism sits at the far end of the spectrum, and getting there requires a property of receptors that most diagrams leave out.

RULE **An antagonist contributes no signal of its own; it only subtracts what is already there.** Everything a pure blocker appears to do is the removal of an existing action, so its apparent strength in a given patient is really a measurement of that patient's agonist tone. Blockade of a quiet system therefore looks like an ineffective drug and blockade of a driven system looks like a powerful one, at identical exposure and identical receptor occupancy.

3.3 Constitutive activity, and the position past antagonism

Receptors are not necessarily silent when unoccupied. Some systems show a low level of spontaneous signal transduction with no agonist bound at all, most visibly where receptor density is high, and that resting output is called **constitutive activity**. Nothing in the standard receptor diagram depicts it, which is why it is routinely forgotten, and it is the hidden variable that makes the last position on the spectrum observable at all. Without it, the silent antagonist and the

inverse agonist would behave identically, because there would be no resting output left for the inverse agonist to remove.

An **inverse agonist** shuts down even that constitutive activity, so the system's output falls below the level it holds when no agonist is present. This is a different act from blockade rather than a stronger version of it: the antagonist restores the resting baseline, while the inverse agonist takes the system underneath it. Set out in order, output runs from the inverse agonist at the bottom, up to the silent antagonist which sits exactly at the no-agonist baseline, then the partial agonist above that, then the full agonist at the top.

The consequence of all this is regional rather than merely theoretical. Where a receptor system carries no constitutive activity, perhaps because receptors are present in low density there, there is nothing extra to switch off and the inverse agonist is indistinguishable from an antagonist. One molecule may therefore act as an inverse agonist in a dense receptor field and as an ordinary blocker in a sparse one, within the same brain, with no change in dose, formulation or timing. Position on the spectrum is a property of the drug and the tissue together, never of the drug alone, and this is the single most useful correction to the way the vocabulary is first taught.

3.4 Drugs labelled antagonist that may not be

This is not a laboratory curiosity confined to receptors nobody prescribes for. Several agents long considered silent antagonists, among them the **5-HT2A antagonists** and the H1 antihistamines, may in some areas of the brain actually turn out to be inverse agonists. Both of those groups are represented across everyday psychiatric prescribing, so this is a reclassification of drugs already in the ward cupboard rather than of compounds confined to a bench.

The honest teaching position here is the uncertain one, and stating the uncertainty is what earns the mark. The clinical differentiation between an antagonist and an inverse agonist remains to be clarified, so a candidate who asserts a settled clinical difference is over-claiming and can be pushed off it by one follow-up question. What can be said firmly is the mechanism-level statement: the two labels diverge only where constitutive activity exists, they converge exactly where it does not, and receptor-binding affinity alone never decides which of the two a drug is in a particular region, because affinity describes how tightly a ligand sits rather than what it does once seated.

TRAP **Treating the word antagonist in a drug's product literature as a pharmacological finding.** It is a classification, frequently assigned from work in a preparation that had no constitutive activity to reveal. Candidates read the label, conclude the drug can only subtract, and are then unable to explain why a supposedly neutral blocker produced an effect in a system where there was little agonist tone to remove. Answer at the level of the spectrum, name constitutive activity as the deciding variable, and say plainly that the clinical distinction is not yet settled.

3.5 The partial agonist's dual personality

The position between full agonism and silent antagonism is the one that matters most at the moment of prescribing, because it is the only position whose net effect changes sign depending on the patient. In the absence of a full agonist, a **partial agonist** increases signal transduction. In

the presence of a full agonist, the very same molecule at the very same receptor turns the strength of downstream signalling down, because it now occupies receptors that the more powerful natural ligand would otherwise have driven harder. For that reason partial agonists are also called **stabilisers**, a word that describes what they do to a system rather than what they do to a receptor.

Read as a prescribing statement, this is a genuinely strange property. Direction of effect is not a fixed attribute of the drug. It is the drug crossed with the prevailing agonist tone, and tone differs between patients, between phases of illness in the same patient, and between regions within one brain at one moment. A single administration of a partial agonist can therefore be raising transmission in one place and lowering it in another simultaneously. No dose adjustment repairs that, because the variable being crossed with the drug is not dose.

POSITION ON THE SPECTRUM	GIVEN WHERE AGONIST TONE IS LOW	GIVEN WHERE A FULL AGONIST IS ABUNDANT	EFFECT ON RESTING CONSTITUTIVE OUTPUT
Full agonist	Drives signalling to the greatest extent the receptor permits	Adds little; the receptor is already being driven maximally	Raised, along with everything else the receptor does
Partial agonist	Net agonist: output rises toward the molecule's own set point	Net antagonist: output falls toward that same set point	Raised toward the set point, not beyond it
Silent antagonist	No change at all; there is nothing present to block	Blocks the agonist and contributes nothing of its own	Unchanged, which is part of the definition
Inverse agonist	Output falls below the unoccupied baseline	Blocks the agonist, then suppresses the system further	Suppressed; this is its defining action

Two columns of that table describe the same drug, and the reason candidates find partial agonism slippery is that they try to hold one answer for it. There is no one answer. The clean way to carry it is as a conditional: agonist where the system is quiet, antagonist where the system is loud, and the crossover sits at the molecule's own set point.

3.6 The ceiling, and why the dose knob stops working

The second half of partial agonism is the half that changes what a prescriber does on Monday morning. Each partial agonist carries its own **set point** engineered into the molecule, so that no increase in dose raises its receptor activation beyond that point. This ceiling is not a safety limit imposed by a regulator, and it is not a tolerability limit imposed by adverse effects. It is a property of the chemistry of the ligand, it is reached whether or not the prescriber intends it, and it cannot be titrated away.

Three prescribing consequences follow. The first is that escalating a partial agonist because the response is inadequate is, at that receptor, futile: whatever additional effect appears at the higher dose must be arising somewhere other than the receptor whose ceiling has already been reached, which in practice means from the molecule's other actions. The second is that the drug is unusually predictable at the top of its range and unusually patient-dependent at the bottom,

which is the reverse of the intuition trained by ordinary dose-response reasoning. The third is that a partial agonist and a full agonist at the same receptor cannot be compared by dose at all, because they are not on the same output scale; a milligram-for-milligram comparison between them is meaningless however carefully it is calculated. The general relationship between dose and effect, and the therapeutic index it generates, are taken up elsewhere in these notes.

PITFALL

Titrating a partial agonist upward because the response has been partial. The ceiling belongs to the molecule, so once the set point is reached the additional milligrams buy adverse-effect exposure and no further receptor activation. The corrective is to change the mechanism rather than the number: a partial response to a partial agonist is a reason to reconsider which position on the spectrum this patient needs, not a reason to keep climbing a scale that has already ended.

3.7 A worked exemplar, and the dopamine partial agonists

The cleanest worked illustration of the two-sided behaviour comes from outside psychiatry's usual receptor set. Acting as a partial agonist at **alpha-4 beta-2** nicotinic acetylcholine receptors, varenicline activates those receptors to a lesser extent than the full agonist nicotine does, and at the same time prevents nicotine from binding to them. Those are not two separate pharmacological effects that happen to coexist; together they are the definition of the class. Something is delivered, and the stronger natural ligand is kept out. The clinical use of that molecule in smoking cessation belongs to the Substance use disorders notes.

Transpose the identical logic onto dopamine and the D2 partial agonists follow, of which aripiprazole and cariprazine are the familiar examples. Where dopaminergic tone is high, such an agent behaves as a functional antagonist and brings signalling down toward its set point. Where tone is low, the same agent supplies a floor of activation rather than removing what little activity is there. This is why these drugs behave unlike a pure blocker in a way that is visible clinically and is not explained by receptor affinity, and it is also why their behaviour cannot be predicted from a binding table alone. The pathway-by-pathway consequences of dopamine occupancy, and how these agents sit against the rest of their class, are taken up elsewhere in these notes; choosing between them for an individual patient belongs to the Schizophrenia: management, antipsychotics and adverse effects notes.

MNEMONIC

Full, Partial, Silent, Inverse

The spectrum read downwards by signal transduction output: **F**ull agonist at the top, **P**artial agonist below it, **S**ilent antagonist resting exactly at the no-agonist baseline, **I**nverse agonist beneath that baseline. Recited in that order, the antagonist lands where it actually belongs, in the middle of the scale, instead of at the bottom where the everyday word blocker tempts candidates to put it.

GOOD TO REMEMBER

Before predicting what any receptor drug will do to a particular patient, ask two questions in that order: where does this molecule sit on the spectrum, and how much natural transmitter is likely to be at that receptor in this clinical state. A blocker given into a quiet system and a partial agonist given into a driven system both look like inert drugs, for opposite reasons, and neither situation is the kind of treatment failure that a higher dose repairs.

A partial agonist is a net agonist where transmitter tone is low and a net antagonist where tone is high, and no increase in dose will carry it past the set point built into the molecule.

4 · Pharmacokinetics: half-life, steady state and the decisions that follow

Half-life is the parameter that answers the questions a prescriber actually asks. When may I judge whether this dose is working. When may I send a level and believe it. How long after the last tablet the drug is still doing something. Whether a missed dose has cost anything worth replacing. Each is a kinetic question wearing clinical clothes, settled by one number and its arithmetic rather than by the drug's receptor profile. That is why this material earns marks out of proportion to its length: it is generic, it transfers unchanged to every class in these notes, and it can be reconstructed rather than memorised.

What follows treats the kinetic parameters as prescribing instruments and works forward from each one to the decision it settles. The algebraic derivation of clearance and volume of distribution, and the absorption, distribution, metabolism and excretion cascade taught as biochemistry, belong to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes. Which isoenzyme clears which psychotropic has its own section here. Target plasma concentrations, the therapeutic drug monitoring framework built on them, and prescribing where an organ has failed belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

Four to five half-lives	Three half-lives	One or two hours	90 to over 99 per cent
TO A STEADY PLASMA CONCENTRATION	ENOUGH FOR A WORKING APPROXIMATION OF IT	THE MOST A DOSE MAY BE HELD FOR A TROUGH SAMPLE	BOUND FRACTION ACROSS THE ANTIPSYCHOTICS

4.1 The four parameters, and the decision each one governs

Drug disposition is described by four parameters, and the reason to keep them apart is that each answers a different question at the bedside. How much of what was swallowed arrives is one question; where it goes is a second; how efficiently the body removes it is a third; how quickly the circulating concentration falls is a fourth. Conflate any two and the clinical inference fails, because a drug can be widely distributed and rapidly cleared at once.

PARAMETER	WHAT IT MEASURES	THE DECISION IT GOVERNS
Bioavailability	The fraction of an administered dose absorbed into the systemic circulation	Whether a change of route or formulation changes exposure, and how much of a swallowed dose is worth counting
Volume of distribution	The apparent space in the body available to contain the drug	How large the tissue reservoir is, and therefore how long the drug keeps returning to plasma once dosing stops
Clearance	The efficiency with which the body eliminates drug from the systemic circulation	How fast exposure falls, and what an inducer, an inhibitor or any change in metabolic capacity will do to it
Elimination half-life	The rate at which drug is removed from the systemic circulation	When steady state arrives, when a level means anything, and what abrupt stopping will feel like

The same parameters split along the time axis, and that split is what most stems test. The front of the sequence sets onset, the back sets duration: absorption and distribution decide how quickly an effect appears, while metabolism and excretion end that effect by clearing the active form out of the body. A patient who says the hypnotic takes too long to work is describing an absorption problem; one who says the same hypnotic leaves them heavy next morning is describing an elimination problem. The complaints sound alike, sit at opposite ends of one sequence, and are fixed differently.

MNEMONIC

In, Around, Out, Over time

In is bioavailability, how much of the dose arrives. **Around** is the volume of distribution, the space it spreads into. **Out** is clearance, the efficiency of removal. **Over time** is the elimination half-life, the rate at which what is circulating disappears. Four words, four parameters, in the order the drug meets them.

4.2 First-order elimination, and the drugs that break it

Almost everything prescribed in psychiatry obeys one elimination rule. **First-order kinetics** means that clearance stays constant when it is expressed as a volume per unit time, so the mass of drug removed rises as the circulating concentration rises. The body is not working harder at high concentrations; it cleans the same volume of blood in the same time, and that volume simply contains more drug. Two consequences follow and both are examinable. A constant fraction of what is present leaves in each interval, which is the only reason a half-life is a stable number at all; and a dose change produces a proportionate change in the eventual concentration, which makes ordinary titration predictable.

The exception is what happens when the machinery runs out of capacity. **Zero-order kinetics**, also written as saturation or non-linear kinetics, applies once the metabolising or eliminating

mechanisms are saturated, so a fixed amount of drug leaves per unit time whatever the plasma level happens to be. Past that point the orderly relationship between dose and concentration collapses, and a further increment in dose produces a disproportionate rise in concentration. The half-life also stops being a fixed property of the molecule, because the time taken to halve a concentration now depends on how high it was to begin with. This is the pharmacology behind toxicity that appears from nowhere after a modest dose increase.

FEATURE	FIRST-ORDER ELIMINATION	ZERO-ORDER ELIMINATION
What stays constant	Clearance, as a volume per unit time	The mass of drug removed per unit time
What rises with concentration	The mass of drug removed	Nothing: removal is already at capacity
Effect of raising the dose	A proportionate rise in concentration	A disproportionate rise in concentration
Status of the half-life	A stable property of the drug	Not fixed: it depends on where the concentration started
Where it is met	Almost every psychotropic at usual doses	A short list, and the sources do not agree on it

TRAP **Answering "name the zero-order drugs" with the psychotropic list, or answering "which psychotropic behaves unpredictably on dose escalation" with the classical trio.** The standard references print two lists because they answer two questions, and neither replaces the other. The classical examination answer is ethanol, phenytoin and high-dose salicylate, phenytoin being the one that also appears in psychiatric prescribing tables. The psychotropic answer is different: paroxetine and nefazodone are described as showing zero-order kinetics at clinically relevant doses, with fluoxetine and fluvoxamine at high dose reported as non-linear. The mark is lost by producing the memorised trio for a stem describing a patient whose antidepressant concentration jumped after a modest increase. If the question is open, give both and say why they differ.

4.3 Half-life is a derived number, not a fixed property

Stated formally, the elimination half-life is the time required for the concentration of drug to fall by one half. Its real importance, and the point most often missed, is that it is not a primary property of the molecule at all. Half-life follows mathematically from hepatic clearance and from the volume of distribution, so it lengthens whenever distribution widens or clearance falls. Learn the direction rather than the formula: anything that enlarges the space the drug occupies, or slows the machinery removing it, buys a longer half-life, and that moves every downstream decision at once. The equation stating the relationship formally belongs to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

Two readings follow at the bedside. The first is that a half-life printed in a formulary is a population figure from reasonably healthy volunteers, and the patient in front of you sits wherever their own clearance and distribution put them: treat it as a starting estimate whose direction of error is predictable. The second is that the half-life, not the hour, is the natural unit

of time here, because the same count of half-lives means the same thing for a fast-cleared drug and a slow one, which is what makes the rule below portable across every class in these notes.

4.4 Steady state, and the arithmetic that settles an old argument

A plasma concentration means what a textbook says it means only at one point in treatment. **Steady state** is the equilibrium at which the amount of drug going in equals the amount coming out, so there is no net change in concentration across successive doses, and it is reached after **four to five half-lives**. That same interval does three jobs, the single most useful thing here: it is also the time needed to reach a new steady state after the daily dose has been raised or lowered, and the time needed for complete elimination once the drug is stopped altogether.

Accumulation towards the plateau is logarithmic rather than linear, and writing the curve out settles a disagreement between standard texts before it can start. Two half-lives reach 75 per cent of steady state, three reach 87.5 per cent, four reach 94 per cent and five reach 97 per cent. Each interval closes half of whatever gap remains, so the concentration climbs steeply and then crawls. Almost everything clinically interesting has happened well before the plateau is formally complete, which is why the last stretch of the curve is where textbooks stop agreeing.

RULE **The dose sets the height of the plateau; the half-life sets the time taken to reach it.** Nothing done to the maintenance dose brings steady state closer, because raising the dose raises the target as fast as it raises the concentration. That one sentence generates the loading-dose rule, the level-timing rule and the washout rule, with no further arithmetic.

TRAP **Treating three half-lives and four to five half-lives as rival answers to be chosen between.** Both are printed in standard references, and a candidate who has met only one assumes the other is an error. The Maudsley Prescribing Guidelines in Psychiatry, 15th edition accepts the familiar adage and then qualifies it, stating that **three half-lives** are already sufficient for an approximation of steady-state concentrations, while other standard texts print the longer figure without qualification. The accumulation percentages show this is not a disagreement about fact but about tolerance: one curve read to a looser standard and a tighter one. Answer with the longer figure as the working rule, then name the qualification. Naming it shows the curve has been understood rather than the phrase memorised.

4.5 The four decisions that steady state governs

Four ordinary decisions turn on the arithmetic above, and each goes wrong in a recognisable way when it is skipped.

- **When to judge a dose.** A dose that has not plateaued has not shown what it can do. Escalating early reads the upslope of the accumulation curve as treatment failure, and the patient ends on a higher dose than the illness required.
- **Whether to load.** **A loading dose** is used to reach therapeutic concentrations as quickly as possible, and it does exactly that. What it does not do is bring steady state forward. The distinction is not academic: loading buys an early effect and nothing else.
- **When to take a level.** A concentration measured before the plateau is a snapshot of an accumulating system and cannot be compared with any published figure, however well drawn.

Wait for it, and wait again after any dose change, for the same interval.

- **How to take it.** **A trough sample** is drawn immediately before the next dose is due, and **the next dose must not be withheld for more than one or possibly two hours** in order to fit the sampling in, because a longer delay returns a falsely low result.

PITFALL

Reacting to a low level drawn too early. The number is real, the assay is right, and the inference is still wrong, because a concentration taken before the plateau measures where the patient has got to, not where they will end up. The dose is increased, the plateau then arrives on top of the increase, and the patient is over-exposed for reasons unrelated to their metabolism. The same logic runs in reverse after a reduction. What the numbers should read, and which drugs are worth monitoring at all, belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

IN INDIA

Where a level is drawn opportunistically at whatever hour the patient reaches a busy outpatient department rather than at a scheduled pre-dose time, two decisions stay entirely within the clinician's control. First, do not send the sample until the plateau has arrived: an early result cannot be interpreted against any reference figure and usually prompts an unnecessary increase. Second, if the sample is meant as a trough, take it immediately before the next dose is due and do not hold that dose back beyond one or possibly two hours for convenience, because a longer wait returns a falsely low value. Where the timing cannot be arranged, record the interval since the last dose on the request form and interpret the result against it, rather than treating an untimed number as a trough.

4.6 Distribution: bound drug, free drug and the tissue reservoir

Two features of distribution explain most of what is otherwise puzzling here. The first: the concentration the laboratory reports is not the concentration doing the work. Only **free drug** is available for activity at the receptor, and psychotropics are heavily protein bound; taking the antipsychotics as the worked example of the class, **the bound fraction runs from around 90 per cent to more than 99 per cent**. A total concentration is therefore a proxy, good only while the bound fraction stays stable. Anything altering binding alters the free fraction without necessarily altering the total, which is why interpreting levels in physical illness is treated separately in the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

The second feature is lipophilicity, the same property that makes these drugs work at all. Lipid solubility lets a psychotropic cross the blood-brain barrier, and it also distributes the drug into peripheral tissue on a large scale. The apparent volume of distribution is correspondingly large, and drug held in that peripheral store seeps back into the circulation long after dosing has stopped. This closes a loop: wide distribution lengthens the half-life directly, and it means plasma concentration alone understates how long the body still holds the drug. Efficacy, adverse effects and active moieties can all outlast the last dose.

4.7 Getting in: bioavailability and the first pass

Bioavailability is where the arithmetic starts, and for oral drugs the dominant influence on it is the liver. A drug absorbed from the gut reaches the portal circulation before anywhere else, and **presystemic extraction**, the first-pass effect, removes a proportion of it during that passage. For the antipsychotics the consequence is stated bluntly in the standard reference: the bulk of an orally administered antipsychotic dose does not reach the systemic circulation at all. Scope that claim carefully: it is made of the antipsychotics as a class, not of every oral psychotropic.

A parenteral route bypasses the portal circulation, so the same milligram figure by a different route is not the same exposure, and dose equivalence across routes is never arithmetic alone. The release kinetics of long-acting injectables, where absorption becomes rate-limiting and the apparent half-life is set by the formulation rather than the molecule, have their own section here. And because first-pass extraction is hepatic, anything changing enzyme activity changes bioavailability as well as clearance, which is where the substrate map taught elsewhere in these notes joins this one.

4.8 Stopping: half-life as the discontinuation variable

The last decision half-life governs is the one made when treatment ends, and it is the one most often made carelessly. The more abruptly a drug is stopped and the shorter its elimination half-life, the more likely clinically significant withdrawal symptoms become. Both variables are at least partly under the prescriber's control: the speed of the reduction always, and the choice of agent at the outset. A short half-life gives a steep fall between doses and a steeper one after the last, so systems that adapted to the drug get no time to readapt. A long half-life provides a built-in taper, because elimination is itself slow.

Two corollaries are worth carrying. The washout interval is the one already met: complete elimination takes the same count of half-lives as reaching steady state did, so a drug stopped now is still present, in falling amounts, throughout. And a missed dose is a miniature version of the same event, which is why it matters far more for short-half-life agents, and why adherence and discontinuation are pharmacologically one problem at different scales.

> **GOOD TO REMEMBER**
>
> Before stopping or switching anything, ask two questions in this order: what is the half-life, and how abrupt is the stop. Short and abrupt, and the taper is not optional and the patient should be told what to expect. Long and abrupt, and the drug tapers itself, so a patient distressed well after the stop is more likely relapsing than withdrawing, a different problem with a different answer. The clinical management of antidepressant discontinuation belongs to the Mood disorders: pharmacotherapy notes; benzodiazepine withdrawal as a clinical syndrome belongs to the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes, and as a substance use disorder to the Substance use disorders notes.

Time to steady state, time to a new steady state after a dose change, and time to complete washout after stopping are one and the same interval, counted in half-lives and set by the half-life alone.

5 · Cytochrome P450: the psychotropic substrate map in practice

A cytochrome table is not a list to be memorised, it is a prediction engine. Two patients on the same milligram of the same drug can sit an order of magnitude apart in plasma concentration, and once dose, adherence and formulation are excluded, most of what remains is enzyme. What follows treats the substrate map as a working instrument: which isoenzymes carry the load, what inherited variation does to them, how inhibition and induction differ in speed as well as direction, and where the map stops predicting. The enzymology underneath it, from the nomenclature and haem mono-oxygenase chemistry to the phase I and phase II division, belongs to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes. The cross-class interaction matrix and therapeutic drug monitoring with its target ranges belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

It is examinable because it generates decisions reachable no other way: why a level collapses on discharge from a smoke-free ward, why an unremarkable dose overshoots in another population, and why some interactions declare themselves at once and others take a working week.

60% OF P450 CONTENT, THE CYP3A SHARE	**About 20%** CYP2C19 POOR METABOLISERS IN ASIAN POPULATIONS	**Up to 50%** FALL IN CLOZAPINE AND OLANZAPINE LEVELS IN SMOKERS	**Every 2 days** HALVING OF ACTIVITY ONCE SMOKING STOPS

5.1 What the map buys, and where prediction stops

The case for the map is modest. **Knowing which enzyme handles which drug** lets a clinician predict or confirm interactions that neither the regulatory trial programme nor ordinary clinical use had uncovered, and inherited variation in those enzymes explains much of the otherwise unaccountable variation between patients. A trial programme tests the combinations somebody thought to test; your patient is on a combination nobody tested.

Prediction is not determination, and the source offering the map says so. The effects of polymorphism and of pharmacokinetic interaction are difficult to predict because many drugs are handled by more than one enzyme, so an **alternative pathway** may compensate when one route is inhibited. A drug with two competent routes of clearance is a road with a diversion: block one and traffic still moves. That is why the same inhibitor is irrelevant for one substrate and decisive for another. The second limit is anatomical: the cytochromes are active at sites other than the liver, the gut and the brain among them, so part of a drug's exposure is settled before it reaches a hepatocyte.

RULE **Name the enzyme before you name the interaction, then ask whether that enzyme is the only route.** An interaction predicted from a table is a hypothesis about one pathway; whether it reaches the patient depends on how much clearance travels that pathway and whether anything else takes up the slack. The map tells you where to look, not how much you will find.

5.2 The isoenzymes that carry the load, and their unequal share

Only a handful of isoenzymes matter for psychotropic prescribing, and they are not equal partners. In the comparative enzyme table of the Maudsley Prescribing Guidelines, the CYP3A column is given as 60 percent of P450 content, the largest share of any enzyme in the table. Two cautions travel with it. It is tabulated for **the CYP3A subfamily rather than for CYP3A4 alone**, so it is not a statement about a single gene product, and the cell says P450 content without specifying that the content is hepatic. It is also tempting, and unsupported, to conclude from abundance alone that this is why CYP3A4 carries the most psychotropic substrates: abundance and substrate breadth are separate properties.

What matters more than abundance is what varies. Some isoenzymes vary chiefly by genotype, some by exposure, and one entry below is not an enzyme at all.

ISOENZYME OR TRANSPORTER	WHAT THE SOURCE STATES ABOUT ITS VARIATION	CONSEQUENCE AT THE POINT OF PRESCRIBING
CYP1A2	Induced by the polycyclic aromatic hydrocarbons in tobacco smoke, this enzyme in particular	Substrate levels sit low in a smoker and climb when smoking stops
CYP2C19	Also induced by smoking; poor metabolisers about 20 percent of Asian people against 3 to 5 percent of people of European descent	Genotype, not only co-prescription, sets the exposure a standard dose produces
CYP2D6	No inducer listed in the Maudsley table of psychotropic interactions; poor metabolisers 3 to 5 percent of people of European descent in one table and 7 percent in another	Interactions arise from inhibition or from genotype, never from induction
CYP3A4	The subfamily holds the largest tabulated share of P450 content; possibly induced by smoking	Substrate overlap with P-glycoprotein is common, so two mechanisms move together
UGT	Responsible for phase II conjugation; potently inhibited by valproate	Invisible to a mental map built only from cytochromes
P-glycoprotein	A transporter, not an enzyme, sited in the gut wall, the testes and the blood-brain barrier	Inhibitors are anticipated to raise, and inducers to lower, substrate uptake

5.3 Genotype: four phenotypes, and frequencies the sources dispute

Enzyme activity is under genetic control, and the vocabulary is standard and worth using precisely. Phenotypes are conventionally described as **poor metabolisers**, intermediate metabolisers, normal metabolisers, and rapid or ultrarapid metabolisers. The ladder runs one way, from no useful activity up to more than a standard dose assumes. A poor metaboliser accumulates the drug and presents with dose-related adverse effects on an unremarkable prescription. An ultrarapid metaboliser clears it before it can act and presents as a non-responder, the more dangerous error because the reflex answer to non-response is more dose.

MNEMONIC

Poor, Intermediate, Normal, Rapid

The four metaboliser phenotypes in ascending order of enzyme activity, rapid and ultrarapid at the top. Read the two ends as the presentations they generate: **P**oor looks like toxicity on a normal prescription, **R**apid looks like treatment failure on a normal prescription.

Frequencies are where marks are lost, because the published figures genuinely differ. For CYP2D6 in people of European descent the enzyme table gives poor metabolisers at 3 to 5 percent, while a different table in the same edition of the same book gives 7 percent for European-ancestry North Americans, and the figures quoted from a second textbook where the tricyclics are taught elsewhere in these notes run higher again. None is a misprint: they are different assay panels, reference populations and rules for converting a diplotype into a predicted phenotype.

TRAP **Quoting one poor-metaboliser frequency as though the literature agreed on it.** Candidates memorise whichever number their textbook printed and defend it. The reproducible finding is not the headline percentage but the direction: CYP2D6 poor metabolism is distinctly commoner in people of European descent than in East Asian populations, and CYP2C19 poor metabolism runs the other way. These notes carry more than one figure for the same population, each with its own source. Give the figure with its source and its population, or give the direction.

The CYP2C19 contrast is the one to carry away. **About 20 percent of Asian people are CYP2C19 poor metabolisers, against 3 to 5 percent of people of European descent**, a difference of roughly four to six times. The caveat matters as much as the figure: the source says Asian without disaggregating South Asian, so this is not an Indian prevalence and must never be quoted as one. The same reference tabulates finer ancestry bands, including a South and Central Asian column, and a reader who needs that number should consult the printed table rather than a reproduced one. The CYP2D6 counterpart for people of Asian descent is quoted where the tricyclics are taught elsewhere in these notes.

IN INDIA

Two decisions follow, neither needing a genotype result. For any drug cleared substantially through CYP2C19, treat the standard European-derived starting dose as an upper bound rather than a default, because poor-metaboliser status is several times commoner in Asian than in European-descent populations, so a routine dose produces high exposure in a larger minority of your clinic than licensing studies imply. Titrate against response, not to a target milligram. And when dose-related adverse effects appear at a dose you would call modest, put metaboliser status on the differential beside adherence and interaction, rather than concluding intolerance and abandoning a drug class that might have worked lower down.

5.4 Inhibition and induction: one map, two directions, two clocks

An enzyme can be interfered with in two ways, and the difference is mechanical, which is why it is also temporal. Inhibition blocks a pathway that already exists. Induction increases the quantity of enzyme protein, which requires that the protein first be made and later cease to be made.

Lewis's Child and Adolescent Psychiatry puts numbers on the slow arm: induction takes 3 to 10 days to start and 5 to 12 days to stop, because protein synthesis must first occur and then cease. The fast arm carries a condition that is routinely dropped and should not be. The same source states that **inhibition occurs rapidly for drugs that have a high presystemic clearance**, and that drugs with low presystemic clearance are the more susceptible overall, having more unused metabolic capacity. Four asymmetries follow.

- An inhibitor added to a regimen can raise exposure before the next review; an inducer will still be building its effect days after it was added.
- Stopping an inhibitor unwinds on the timescale of its own elimination; stopping an inducer unwinds only as fast as surplus enzyme protein is degraded.
- A level checked too soon after an inducer is started or stopped measures a system still in transit.
- The two are not symmetrical opposites: an enzyme can be inhibited that cannot be induced.

CYP2D6 is the standing example. The Maudsley Prescribing Guidelines list no known inducer for this enzyme among the psychotropics, so every interaction there is either an inhibition or a matter of inherited phenotype. That is a useful simplification: a fall in exposure at this enzyme is not explained by an inducer, so look instead at adherence, absorption, or the wrong enzyme.

PITFALL

Teaching yourself that inhibition is immediate, full stop. The source conditions the speed on the drug: it is stated for agents with high presystemic clearance, where blocking gut and first-pass metabolism jumps what reaches the circulation. For a drug not extensively cleared before that point, the concentration still climbs to a new steady state at a rate set by its half-life, not by how fast the inhibitor bound the enzyme. Expecting a next-morning level to have moved fully is how a real interaction gets dismissed as absent.

5.5 Tobacco smoke: the induction psychiatry meets most often

The commonest clinically important induction in psychiatric practice is not caused by a drug at all. Tobacco smoke contains **polycyclic aromatic hydrocarbons** that induce hepatic enzymes, CYP1A2 in particular, with CYP2C19 and possibly CYP3A4 and some variants of UGT also induced. Given how heavily smoking clusters with severe mental illness, this is the interaction a psychiatrist meets most.

Its magnitude is not trivial. **Smoking reduces clozapine plasma levels by up to 50 percent, and olanzapine levels by up to 50 percent as well**. Read that as the size of the CYP1A2 induction effect rather than a fact about either agent: clozapine itself is taught in its own section of these notes, and target concentrations belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes. More surprising is how little smoke is needed. The clozapine-lowering effect may be maximal at as few as 2 to 5 cigarettes a day, so the light smoker who denies being one is fully induced, and cutting down is not a proportional reduction.

TRAP **Attributing the interaction to nicotine.** It is the smoke, not the drug in it. Only tobacco smoking induces hepatic enzymes in this way, cigarettes and cannabis-tobacco joints included; nicotine replacement, vaping devices and electronic cigarettes contain no polycyclic aromatic compounds and have no effect on enzyme activity. The consequence is exact and counterintuitive: a patient switched from cigarettes to a patch or to vaping has stopped inducing, so levels rise as though they had quit outright, while their nicotine intake is unchanged and they do not consider themselves to have stopped.

5.6 The transition: what happens when a smoker stops

Cessation is where this topic becomes a prescribing action, because de-induction has its own measurable rate. **When a person stops smoking, enzyme activity halves roughly every two days**. That figure does more work than it looks: the change is neither instantaneous nor slow enough to ignore, most of the movement falls inside the first week, and a level drawn the morning after the last cigarette is obsolete by the time it is reported.

The stop is often not the patient's choice. Admission to a smoke-free ward is an unplanned de-induction imposed on a fixed dose, running in reverse at discharge, and neither transition is written in the notes as an interaction, which is why it is missed.

For a patient established on clozapine for a psychotic illness, the Maudsley Prescribing Guidelines set out a procedure rather than a single number: take a plasma level before the patient stops smoking, then reduce the dose gradually over a week until it sits at around 75 percent of the original, a first reduction of **25 percent**; repeat the plasma level a week after stopping, and anticipate further reductions. Other sources quote a total reduction in the region of 30 to 40 percent, roughly where a staged protocol lands once later adjustments are made, so the two are not in genuine opposition. What is examinable is the shape of the answer: measure, reduce in steps, measure again, expect to reduce further.

Only tobacco smoke induces, nicotine does not, and once the smoke stops enzyme activity halves about every two days, which is why the dose comes down in staged steps against repeated levels rather than in one calculated cut.

GOOD TO REMEMBER

Record smoking status as a number of cigarettes, not a yes or no, at every review of a patient on a CYP1A2 substrate, and again at admission and discharge. A patient reporting a handful a day is already fully induced, so their level is a low induced level and not a true baseline. A planned quit attempt is a pharmacokinetic event and should be dated in advance so a level can precede it. An unplanned change of setting, either direction, should prompt a level rather than a wait for symptoms, because the symptom announcing a rise is usually an adverse effect.

5.7 Beyond the cytochromes: the conjugation arm no CYP table shows

A mental model built only from cytochromes misses an entire family of interactions, and the one it misses is among the most consequential in mood-disorder prescribing. Uridine diphosphate glucuronosyltransferase, abbreviated **UGT**, is the enzyme family responsible for phase II conjugation reactions, and valproate is a potent inhibitor of UGT, which is the basis of its interaction with lamotrigine, a drug cleared primarily by that route. Nothing in a cytochrome table predicts this, because no cytochrome is involved.

The teaching point is the class of error, not the pair of drugs. Where exposure behaves unexpectedly and no cytochrome explains it, conjugation is the next place to look, and the same reasoning covers smoking's effect on UGT variants. The slower lamotrigine titration that follows is taught with lamotrigine elsewhere in these notes, and enzyme induction by an antiepileptic mood stabiliser, autoinduction included, is taught with carbamazepine. The biochemistry of conjugation belongs to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

5.8 P-glycoprotein: the transporter layer above the map

The last layer is not metabolism at all. P-glycoprotein is a drug transporter protein found in the gut wall, and it can actively eject drugs that have passively diffused across that wall; it is also present in the testes and in the blood-brain barrier. Absorption, on this account, balances passive entry against active removal, so a drug can be well absorbed while a transporter returns much of it to the lumen.

The interaction logic mirrors the enzyme logic, and the source's hedge is worth keeping. **Drugs that inhibit P-glycoprotein are anticipated to increase, and drugs that induce it to reduce, the uptake of other drugs that are its substrates.** Because the transporter also sits at the blood-brain barrier, the same mechanism governs how much of a substrate reaches the site of action, a different question from how much reaches plasma. Then comes the observation that makes the map harder to read: many drugs that are substrates for CYP3A4 have also been found to be substrates for P-glycoprotein. Enzyme and transporter are often perturbed by the same

agent in the same direction, so an effect larger than the table predicts may be two mechanisms adding rather than one misjudged, and a plasma level will not separate them.

The four layers answer one question. Which enzyme, in what quantity, with what inherited variation, and is anything at the gut wall acting before the liver sees the drug. Answer those and an unexpected level becomes an inference rather than a mystery. Take the map for a verdict rather than a hypothesis and it fails where it was always going to: on the drug with two clearance routes, in the patient whose ancestry was not the one tabulated, on the morning the ward became smoke-free.

6 · Therapeutic index and dose-response: which psychotropics need a level

Two questions decide how any psychotropic is prescribed, and pharmacology answers both before the patient is seen. The first is how much room lies between the dose that produces the wanted effect and the dose that produces harm. The second follows from it: whether that room is wide enough to prescribe against clinical response alone, or so narrow that a blood concentration has to stand in for a judgement the clinic cannot otherwise make. Most arguments about monitoring, about titration speed, and about which of two equally effective agents to reach for reduce to those two questions. What follows takes them in order: the shape of the dose-response relationship, the ratio built on that shape, and the small set of drugs in psychiatry for which a measured concentration is genuinely informative. The receptor mathematics beneath the curves, the occupancy relations and the molecular basis of tolerance and sensitisation belong to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes; what is developed here is the prescribing consequence that sits on top of them. The published target concentrations themselves, and the monitoring apparatus built around them, belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

1 to 2	More than 100	25 percent	Up to 25 percent
LOWEST THERAPEUTIC INDICES AMONG DRUGS IN USE	THE SAFE END OF THE SAME SCALE	RESPOND BELOW THE PUBLISHED TARGET RANGE	TOLERATE CONCENTRATIONS ABOVE IT

6.1 What pharmacodynamics is asked to deliver

Pharmacodynamics describes what a drug does to the body: the intensity of its effects and the time course over which they appear and fade. By long convention the subject is organised under four headings, and the organisation is worth holding because it maps directly onto four different prescribing decisions. The first heading is the receptor mechanism, which determines what the drug can do at all. The second is the dose-response curve, which determines how the effect grows as the dose grows. The third is the therapeutic index, which determines whether the effect can be pursued safely. The fourth is the development of tolerance, dependence and withdrawal, which determines what happens to the first three with repeated exposure.

Only the middle two are developed here. The receptor mechanisms belong to the section of these notes that reads a receptor profile forward into an adverse-effect signature, and the biochemistry

underneath them to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes. Tolerance, dependence and withdrawal as pharmacological phenomena are taken up in the section on the benzodiazepines, where they are most instructive. The reason for separating the third heading out and giving it a section of its own is practical rather than academic: it is the only one of the four that changes what you do with a syringe and a laboratory form. A drug whose useful and harmful doses are far apart is prescribed by watching the patient; one whose useful and harmful doses are close together is prescribed by watching a number, and the whole apparatus of blood-level monitoring exists because a handful of drugs fall into that second group.

6.2 Two dose-response curves that answer different questions

Candidates lose marks here by treating the dose-response curve as one object. There are two, they are constructed from different observations, and the landmarks read off them mean different things. In a **graded dose-response** relationship, what is measured is the size of a single person's response, and the magnitude of that individual's response usually increases as the dose is increased. In a **quantal dose-response** relationship the response is not measured at all but judged: it is either present or absent in a given individual, and what increases with dose is the percentage of the population that responds.

The distinction is not pedantry. A graded curve is what you are implicitly using when you raise an antipsychotic on the reasoning that this patient is partially better and more will help. A quantal curve is what a trial is reporting when it says that a proportion of patients responded at a given dose, and it says nothing whatever about how much better any one of them became. Sliding between the two is how a population response rate gets misread as a promise to the person in the room.

AXIS OF COMPARISON	GRADED DOSE-RESPONSE	QUANTAL DOSE-RESPONSE
What varies with dose	The magnitude of one individual's response	The percentage of the population that responds
How the response is judged	As a matter of degree	As present or absent in a given individual
Unit of observation	One patient, observed repeatedly	A population, observed once each
Question it can answer	How much more will this patient gain from a higher dose	How many patients gain anything at all at this dose

6.3 The landmark doses, and the ratio built from them

Four abbreviations carry the arithmetic and they are worth learning as a set, because examiners rarely ask for one in isolation. ED50 denotes the **median effective dose**; LD50 denotes the median lethal dose; TD50 denotes the median toxic dose; and TI denotes the therapeutic index. The prefix in each case is a median, which is to say a population statement: the dose at which half the sample shows the event. None of them describes an individual, and that limitation is inherited by everything built on top of them.

The **therapeutic index** is what is built on top of them. In preclinical work it is a literal ratio, obtained by dividing the median lethal dose determined in experimental animals by the median effective dose. Carried across into human prescribing the term keeps its meaning while losing its arithmetic, because no lethal dose is determined experimentally in people. What it expresses in humans is how selective a drug is in producing the desired effect rather than serious adverse effects. Read that way, the index is a statement about margin: how far you can push before the pharmacology that helps becomes the pharmacology that harms.

A word of caution about the formula. Because a lethal endpoint cannot be sought in humans, textbooks often quote a version substituting the median toxic dose for the median lethal one. That substitution is reasonable and widely taught, but it is not stated in the sources relied on here and is not printed as though it were. The defensible examination position is the one given above: a ratio of a harmful dose to an effective dose, measured as lethal over effective in animals and understood as selectivity in humans.

RULE **The therapeutic index is a statement about the distance between two doses, never about the strength of one.** It says nothing about how good the drug is, how fast it works or how much of it you need. It says only how much room the prescriber has to move. Every downstream decision in this section, from titration speed to whether a level is worth taking, is a consequence of how much room the pharmacology has left you.

6.4 How wide the index runs, and what a narrow one obliges

The spread across the pharmacopoeia is enormous. Among agents in current therapeutic use, index values run from as little as **1 to 2** at the dangerous end to **more than 100** at the safe end. An index at the bottom of that range means the dose that treats and the dose that harms are effectively the same dose, and the drug is being given inside its own margin of error. An index at the top means the two are so far apart that ordinary prescribing error cannot bridge them.

Two obligations follow directly, and both are examinable. Drugs with low therapeutic indices must be administered with caution and with therapeutic drug monitoring. Drugs with very high therapeutic indices are considered safe in the absence of a known allergic response in a given patient. Notice how the second is phrased, because the qualification is the whole point: a wide index protects against dose-related toxicity, which is the kind of harm that scales with exposure. It offers no protection at all against an idiosyncratic or immunological reaction, which does not scale with dose and can be produced by a first exposure at any dose. Candidates who read a wide index as a general safety certificate have collapsed two entirely different mechanisms of harm into one.

Therapeutic drug monitoring is therefore not a refinement bolted on to good practice. For the narrow-index drugs it is the substitute for a margin the pharmacology has failed to provide.

6.5 Potency is not efficacy, and the classic worked example

The index is a comparison between doses of one drug. The commonest error in dose-response reasoning is a comparison between doses of two drugs, read as though it were a comparison of their worth. **Potency** is the relative dose required to achieve a given effect, and it is a property of

how efficiently the molecule engages its target. **Efficacy** is a different quantity altogether: the maximum clinical response the drug can achieve.

The worked example every examiner uses is drawn from the antipsychotics, and it is worth reproducing exactly because the numbers are the point. For the same antipsychotic effect, approximately **5 mg of haloperidol** is required where 100 mg of chlorpromazine would be needed, which makes haloperidol the more potent of the two by a considerable margin; the two drugs are nonetheless equal in their clinical efficacy. Potency here is a statement about the number on the prescription and nothing else. It follows that a low milligram dose is not evidence of a strong drug, that comparing two agents by their usual daily dose compares nothing of clinical interest, and that dose equivalence tables are conversion tools rather than rankings. Which agent to choose within the class, and on what pharmacological axes the agents differ, is developed in its own section of these notes; clinical selection for a patient with schizophrenia belongs to the Schizophrenia: management, antipsychotics and adverse effects notes.

TRAP **Reading greater potency as greater efficacy.** The stem gives two agents with widely different milligram doses and asks which is the stronger drug. Candidates answer with the smaller number, because everyday language uses strength and potency interchangeably and because a small dose feels concentrated. The haloperidol and chlorpromazine pairing is set precisely because it is a large potency difference sitting on top of equal efficacy. If a question offers a milligram comparison and asks about effectiveness, the comparison is a distractor.

6.6 Where the psychotropics sit, and why it shaped the pharmacopoeia

Taken as a group, the psychotropics sit at the comfortable end of the scale, and this is the single most useful generalisation in the section. For most psychotropic drugs the therapeutic index is high. That is why psychiatric prescribing is normally conducted by titrating against clinical response and adverse effect rather than against a laboratory value, and why routine plasma levels are not part of ordinary antidepressant or antipsychotic practice. The exception is named in the same breath by every standard text: where the therapeutic index is low, as it is for lithium, careful monitoring of serum levels is essential. Lithium as the archetypal narrow-index drug, and what monitoring actually buys, has its own section in these notes; its clinical use in bipolar disorder belongs to the Mood disorders: pharmacotherapy notes.

The index is not only a monitoring trigger. It has functioned historically as an argument for whole classes, and the clearest case in psychiatry is the displacement of the tricyclics. The wide therapeutic index of the SSRIs was a very important consideration in their adoption as first-line treatments. The consequence was measurable at population level: deaths from overdose of tricyclic derivatives fell steeply once the SSRIs supplanted that class as the mainstay of treatment. This is the argument to have ready when asked why a class with no clear advantage in efficacy became first-line. Two agents of equal efficacy are not equal drugs if one is far more dangerous in overdose. The class pharmacology of the tricyclics, and the receptor promiscuity that generates their toxicity, is developed in its own section here; overdose management belongs to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

GOOD TO REMEMBER

The therapeutic index is a prescribing variable, not a viva term. When two agents are genuinely comparable in efficacy for the indication in front of you, and the patient has unsupervised access to the supply, the index is a legitimate tie-breaker in its own right. That reasoning reorganised the antidepressant pharmacopoeia, and it applies to an individual prescription as much as to a formulary.

6.7 Which psychotropics actually need a level

Which drugs need a level has a principled answer rather than a customary one. A drug earns a level when four conditions hold together, and the conditions are worth reciting because each of them independently removes an alternative way of getting the dose right.

- A narrow therapeutic index, so that the margin for error is small enough that guessing is unsafe.
- Significant consequences associated with toxicity, so that being wrong matters more than the cost and discomfort of the test.
- Wide interpatient variability, so that the same dose in two people does not produce the same exposure and the dose alone cannot be trusted.
- No clinical endpoint to guide dosing, so that watching the patient does not tell you whether you have arrived.

A fifth requirement is usually stated alongside these and is easy to forget: the drug concentration must bear a proportional relation to the pharmacologic action at the receptor site. Without it the assay measures something irrelevant, however precisely it measures it. This is the condition that quietly disqualifies most psychotropics from routine monitoring, and it is the reason a level is not simply a good idea for every drug that happens to be measurable.

Applied to psychiatric practice the list that satisfies the whole set is short: lithium, valproate, carbamazepine and nortriptyline are the named candidates for therapeutic drug monitoring. Each of these agents has its own section in these notes; the target concentrations themselves, and the framework for acting on a result, belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

MNEMONIC

Narrow, Nasty, Variable, Blind

The four conditions that make a drug a monitoring candidate: a **N**arrow therapeutic index, **N**asty consequences if toxicity occurs, wide between-patient **V**ariability in exposure, and a clinical picture that is **B**lind, offering no endpoint to titrate against. Remove any one and the level stops earning its place.

6.8 Threshold, window and range: what a published number promises

Once a drug qualifies for monitoring, a second set of concepts governs how the resulting number is read, and they are routinely conflated. The first is the **therapeutic threshold**. For most drugs

the concentration at which therapeutic activity appears is fairly constant across a population, which is what makes a single quoted figure meaningful at all: below it, the drug is largely doing nothing useful in most people. A threshold is a floor, and a floor is a comparatively robust thing to state.

A **therapeutic window** is a more ambitious object, because it claims a ceiling as well as a floor. A window defined for a population expresses a range of concentrations for which the likelihood of efficacy is high and the probability of adverse effects is low, but it does not guarantee efficacy or safety for any single individual. That final clause is the teaching point, and it is the part candidates omit. The word range in the phrase therapeutic range describes a distribution, not a rule, and one standard paediatric reference goes further, calling that phrase a misnomer of sorts, since even when the drug concentration lies within the range not every patient will respond and a percentage will experience toxicity.

The size of that mismatch is quantified, and the figures are more startling than most clinicians expect. **Around 25 percent of patients respond below the target range**, and up to 25 percent tolerate blood concentrations above the target range. Put plainly, a published range is constructed so that it captures the middle of a population, and a substantial minority sit outside it in both directions while being perfectly well managed. A patient who is well below the range does not need a dose increase to reach it. A patient who is above it and entirely asymptomatic does not automatically need a reduction. What the range tells you is where the probabilities lie, which is what a population statistic can offer and no more.

PITFALL

Titrating a patient who is already well up into the published range because the level came back below it. The number was obtained to answer a question, and if the patient is in remission the question has already been answered by the patient. Chasing a range in an asymptomatic responder converts a stable prescription into a toxicity risk, and it happens most often when the level is ordered reflexively rather than to settle a specific doubt about adherence, absorption or an interaction.

Most psychotropics have a high therapeutic index and are titrated against the patient; where the index is narrow, as with lithium, a concentration replaces the clinical endpoint that is missing, and even then the published range is a population statement that substantial minorities sit outside in both directions.

Antipsychotics

7 · What typical and atypical actually predict, and what the labels hide

Four separate classifications of the antipsychotics are in simultaneous use, and only one of them names a mechanism. One sorts the drugs by chemical structure. One splits them into typical and atypical. One calls them first and second generation. The fourth sorts them by

pharmacological activity. These are not four descriptions of a single division. They were built at different times, from different evidence and for different purposes, and they disagree both about which agents belong together and about what membership entails. Examination questions cluster here because the labels are so familiar that few candidates have ever asked what any of them encodes.

What follows reads the classification as a pharmacological object: what each system was built from, what it therefore predicts, and the point at which it stops predicting. Which agent to select for a particular patient, how to sequence trials and what to do when they fail belong to the Schizophrenia: management, antipsychotics and adverse effects notes. The reason to be exact is not taxonomic tidiness. Two of these systems are read as though they carried information they do not carry, and the errors that follow are clinical: an agent chosen because it is called atypical, and expected on that basis to leave prolactin alone.

1959	1971	1990
CLOZAPINE SYNTHESISED	OLANZAPINE FIRST PATENTED	THE CUT POINT THAT CREATES THE SECOND GENERATION

7.1 The chemical taxonomy, and what it was built from

The original scheme was frankly structural. Before the 1990s the antipsychotics were sorted by their chemistry, and the sorting began with the first agent of all. Chlorpromazine is a **phenothiazine**, and once that family had been established further phenothiazines were generated and marketed alongside the chemically similar **thioxanthenes**. What came next was different in kind. Entirely separate chemical structures were developed according to pharmacological reasoning rather than by working outwards from an existing ring system, and they produced the **butyrophenones**, the **diphenylbutylpiperidines** and the **substituted benzamides**.

Two features of that history are worth holding on to. The first is that the taxonomy is genuinely chemical: membership is decided by the molecule, not by what the molecule does to a patient, so a family can contain agents of very different potency. The second is that the later families were designed rather than stumbled upon, which is why the structural divisions do not run parallel to any single pharmacological axis. A chemical family is a statement about synthesis, and any predictive power it turns out to have is an empirical bonus rather than a definition. The historical accident worth noticing is that this scheme, the least fashionable of the four, is the only one of them that predicts an outcome it was not built from.

7.2 Structure as a predictor of movement liability

That prediction concerns the extrapyramidal burden, and it is ordered. Three families sit at the top of the ordering among the older drugs: the **piperazine phenothiazines**, together with the butyrophenones and the thioxanthenes. Two sit at the bottom, the **piperidine phenothiazines** and the benzamides. The remaining families, the aliphatic phenothiazines and the diphenylbutylpiperidines, are placed by the source somewhere in between, and placed tentatively: this ordering is firmest at its ends.

STRUCTURAL FAMILY	AGENTS THE SOURCE NAMES	EXTRAPYRAMIDAL LIABILITY
Piperazine phenothiazines	Fluphenazine, trifluoperazine	Most likely of the older drugs
Butyrophenones	Haloperidol	Most likely of the older drugs
Thioxanthenes	Flupentixol	Most likely of the older drugs
Aliphatic phenothiazines	Chlorpromazine	Intermediate, stated tentatively
Diphenylbutylpiperidines	Pimozide	Intermediate, stated tentatively
Piperidine phenothiazines	Pipotiazine	Least likely of the older drugs
Substituted benzamides	Sulpiride, amisulpride	Least likely of the older drugs

Read the table as a gradient rather than as three boxes. It says that a prescriber who knows only a molecule's ring system already knows something real about the movement disorder that molecule tends to generate, which is more than any other label here delivers. The phenomenology of those movement disorders, from acute dystonia through to the tardive syndromes, is taught elsewhere in these notes; what belongs here is only the ordering.

MNEMONIC

PBT at the top, PB at the foot

Piperazine phenothiazines, **B**utyrophenones and **T**hioxanthenes carry the highest extrapyramidal liability of the older agents; **P**iperidine phenothiazines and **B**enzamides carry the lowest. Both ends of the ordering begin with P and B, which is exactly the point: piperazine and piperidine differ by a syllable and sit at opposite ends of the same ordering, so the family name has to be read whole and never by its first half.

PITFALL

Carrying the structural rule forward into the modern agents. The ordering above is stated for the older drugs only, and the same source records that structure-activity relationships of this kind are absent for the newer ones. Reasoning from a modern molecule's chemical family to its movement liability is therefore not conservative practice, it is guesswork wearing the costume of pharmacology. Beyond the older families the question has to be answered from receptor profile and from observed liability, never from the ring system.

7.3 Where the word atypical came from

The second taxonomy was built from a single adverse effect. What originally drove the division into typical and atypical was relative liability to induce extrapyramidal side effects, which means the split was never a statement about mechanism, about receptor occupancy or about efficacy. It was a statement about one column of the side-effect table.

The index case makes this unusually clear. Clozapine had long been called **atypical** on two grounds: a low liability to cause extrapyramidal effects, and a failure in animal-based antipsychotic screening tests. Both grounds are negative. The drug earned its label by not doing things, once in a patient and once in a screening paradigm, and a category defined by absence has

no positive content to transmit to its later members. Clozapine's own pharmacology is taught elsewhere in these notes; what matters here is only that the prototype of the atypical class was admitted to it for reasons that say nothing about how it works.

The failure in animal screening is the part candidates skip. Those paradigms were calibrated on the drugs already in hand, so they detected the pharmacology of the existing agents and scored anything unlike them as inactive. A screen built from a class cannot recognise a member that departs from the class. Clozapine failing the test was therefore evidence about the test as much as about the drug, and the word atypical began life as a confession that the available framework did not fit.

7.4 The distinction that survives: how wide the window is

If the split was built from extrapyramidal liability, the honest way to restate it is quantitative rather than categorical. The useful question about any antipsychotic is how wide the band of doses is within which the drug is antipsychotic and not yet parkinsonian, and that band differs enormously between agents that are otherwise described in the same sentence. The band within which haloperidol is antipsychotic without causing extrapyramidal side effects is extremely narrow, quoted at perhaps **4.0 to 4.5 mg/day**; olanzapine keeps its antipsychotic effect across **5 to 40 mg/day** without generally causing those same effects.

Set those two bands side by side and the binary dissolves into a continuum. Nothing categorical separates the drugs. One has a window a prescriber can step over in a single increment, and the other has a window wide enough that most ordinary dose adjustment stays inside it. Note that the source hedges even the wide band: it says only that olanzapine does not generally cause those effects across it, not that its upper end is free of them, and the occupancy account taught elsewhere in these notes places a ceiling that any band eventually meets. Framed this way, atypicality is a property of the dose-response relationship, not a property of the molecule's membership card, and it explains why the same agent can behave typically at one dose and atypically at another. The general pharmacological principle underneath, the size of the interval between an effective exposure and a toxic one, is taught elsewhere in these notes; what the classification adds is that the relevant toxicity here is a movement disorder, so the interval is drawn against that endpoint rather than against a plasma concentration.

RULE **A classification predicts only what it was built from.** The chemical taxonomy was built from structure and predicts extrapyramidal liability among the older drugs. The typical and atypical split was built from extrapyramidal liability alone and predicts that axis alone. The generational split was built from a date, so nothing pharmacological defines it; where the two groups do differ on an outcome, as they do on the movement axis taught elsewhere in these notes, the label is standing proxy for a receptor property rather than acting as a criterion of its own. Every error in this topic comes from asking one of these systems a question it was never constructed to answer.

7.5 First and second generation: a calendar, not a pharmacology

The generational scheme was introduced to escape the word atypical, and it escaped it by saying even less. The rule for membership is a date and nothing more: any agent introduced

since 1990 counts as second generation, and there the definition stops. There is no receptor criterion, no structural criterion, no efficacy criterion and no adverse-effect criterion, because the boundary encodes time of introduction and nothing other than that. Whatever the term **second-generation antipsychotics** is doing in a sentence, it is not stating a pharmacological criterion.

The proof of that is internal to the class, and it is the detail examiners like best. Clozapine, the drug the whole atypical story begins from, was synthesised in 1959, and olanzapine was first patented in 1971; both are nonetheless filed among the second-generation agents. Those are dates of synthesis and of patent, not of marketing, and the classification rests on introduction rather than on either. So two of the most familiar members of the modern group are molecules of considerable age, one of them among the oldest antipsychotics of all, and the label that appears to describe modernity is in fact describing a regulatory and commercial timeline. A candidate who has internalised this will not read the label as a statement about the molecule.

TRAP **Reading the generation label as a statement about the age of the molecule.** Two of the most familiar second-generation agents existed long before the cut point that defines their group, and the group boundary tracks when a drug reached the market, not when it was made. Candidates who have learned the generations as a chronology of discovery answer confidently and wrongly. The label is an administrative fact about introduction, and the safe use of it is as a shorthand for a list of names rather than as a claim about any molecule.

7.6 What the atypical label does not predict

If the generational scheme carries no pharmacological content by construction, the atypical category is worse: it looks as though it should carry some, and it does not. Search for the property all its members share and no candidate holds without qualification. Nothing pharmacological or chemical clearly binds them together as a group on the Maudsley's reading, save one partial exception, a general but not universal preference for D2 receptors outside the striatum. Note the two hedges the source builds into that exception. The preference is general, so it does not hold for every member, and it is a preference rather than a selectivity. Other sources locate the shared feature elsewhere, in the balance of serotonergic and dopaminergic antagonism, which is where the occupancy account in these notes takes it up, and that account states it for most members of the group rather than for all of them. So the two candidate properties carry the same hedge, and neither yields a definition. How occupancy at those receptors translates into both antipsychotic response and adverse effect across the pathways is taught elsewhere in these notes.

Three specific expectations are attached to the word atypical in ordinary conversation. The classification underwrites none of them, and only the first has even a partial answer.

- A shared mechanism. No candidate holds for the whole group: the regional preference is general rather than universal, and the serotonergic property characterises most members rather than all.
- Better efficacy than the older drugs. Membership does not predict improved efficacy over the older agents, with clozapine and a small number of others excepted by the source itself.

- Freedom from hyperprolactinaemia. Membership does not predict the absence of hyperprolactinaemia either, and this is the expectation that misleads most often in practice.

The prolactin point is worth stating in full because it inverts the intuition rather than merely qualifying it. The source names three modern agents, risperidone, paliperidone and amisulpride, as usually worse than typical drugs for prolactin. Three agents that every clinician files under the modern heading are, on this endpoint, worse than the group they were supposed to improve upon. The full within-class ranking of prolactin elevation, weight gain, sedation and movement liability across individual agents is taught elsewhere in these notes; what belongs here is only that the label does not do the ranking for you.

TRAP **Treating atypical as a promise about prolactin.** The category was built from movement liability and from nothing else, so it makes no commitment whatever about the tuberoinfundibular consequences, and three widely used members of it are worse on that endpoint than the older drugs. This is a genuinely different error from the calendar trap: that one misreads what the boundary is drawn on, this one misreads what the boundary implies. Endocrine consequences follow the individual agent's receptor profile, never its group membership.

7.7 The mechanism-based replacement

The obvious remedy for a set of labels that name eras and absences is a set of labels that name actions, and that is what the fourth system attempts. **Neuroscience-based Nomenclature** does exactly that. Each agent is filed under its pharmacological activity rather than its generation, so an antipsychotic stops being first or second and becomes instead a dopamine receptor antagonist, or a **dopamine serotonin antagonist**, according to what it does. The name of the drug class then states the mechanism instead of concealing it, and a reader who has never met a particular agent still learns something true from its category.

SYSTEM	WHAT IT WAS BUILT FROM	WHAT IT THEREFORE PREDICTS
Chemical structure	The molecule's structural family	Relative extrapyramidal liability, among the older drugs only
Typical versus atypical	Relative liability to induce extrapyramidal side effects	That one axis; neither efficacy nor freedom from hyperprolactinaemia
First versus second generation	Time of introduction, and nothing else	Nothing by definition; where the groups differ, the label proxies a receptor property
Neuroscience-based Nomenclature	The agent's pharmacological activity	The mechanism its own name states

The gain is real and it is also limited, and an examination answer that says so is stronger than one that simply endorses the newer scheme. Naming the receptor action does not by itself rank agents on efficacy, does not predict weight change or sedation, and cannot resolve the case of a drug whose clinical behaviour is not explained by the receptors it is named for. What the

mechanism-based names do achieve is the removal of a false implicature: a term built from pharmacology cannot be misread as a promise about outcome, because it never claimed to be one.

7.8 Using the labels without being used by them

Assemble the four systems and a working order of preference falls out. When the agent in question is one of the older drugs, its structural family is the label that carries information, and the extrapyramidal ordering set out above is the thing to recall. When the agent is a modern one, the useful description is its receptor activity and the width of the dose band within which it is effective without producing movement effects; the generational label defines nothing pharmacological and the atypical label points only at the movement axis it was built from. When the question is about efficacy, prolactin or metabolic burden, no label answers it at all and the agent has to be considered individually.

The deeper habit this topic is teaching is scepticism about inherited vocabulary. Each of these systems was reasonable when it was made and was then asked to do work it was never built for, because a convenient name will be used beyond its evidence. That pattern recurs across every other drug class in these notes.

GOOD TO REMEMBER

When a colleague asks whether a drug is typical or atypical, the answerable question underneath is one of two. If the agent is one of the older drugs, ask which structural family it belongs to, because that ordering is real and the piperazine phenothiazines, butyrophenones and thioxanthenes sit at the top of it. If it is a modern agent, ask instead how wide the band of doses is within which it works without producing movement effects, because that is the property the word atypical was always gesturing at. The generation label answers neither question.

The generational split records only when a drug reached the market; the chemical family predicts extrapyramidal liability among the older agents alone; and the only distinction that survives contact with a patient is how wide the dose band is within which an agent is effective without causing extrapyramidal side effects.

8 · Dopamine pathway occupancy: one mechanism, four consequences

One pharmacological action explains most of what this class does to a patient, and it explains the unwanted effects by exactly the same route as the wanted ones. Every antipsychotic in routine use reduces transmission at the dopamine D2 receptor. What separates one agent from another is how much of that receptor population it holds at a clinically used dose, how tightly it holds on, and what else it binds along the way. Occupancy is therefore the organising variable of the class, and it is worth learning as a variable rather than as a list of drug names, because dopamine receptors sit in anatomically separate pathways doing unrelated jobs and a drug given systemically cannot choose between them. One blockade, several destinations, several consequences.

That is the whole logic of this section, and it is what a stem is testing when it asks why an agent that quietens psychotic phenomena also stiffens a patient's gait, or why a young woman develops breast secretion on a drug prescribed for something else. The anatomy itself is neuroscience and belongs to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes; it appears here only as the substrate on which occupancy is spent. The dopamine hypothesis, read as an account of why these agents work in a particular disorder, belongs to the Schizophrenia: management, antipsychotics and adverse effects notes.

60 to 70 per cent	more than 90 per cent	about 40 per cent	around 80 per cent
STRIATAL D2 OCCUPANCY A PLAIN D2 ANTAGONIST NEEDS	OCCUPANCY A D2 PARTIAL AGONIST NEEDS	ENOUGH FOR CLOZAPINE, QUETIAPINE AND LUMATEPERONE	CEILING ABOVE WHICH EXTRAPYRAMIDAL EFFECTS APPEAR

8.1 Occupancy is the quantity that travels; dose is not

A dose in milligrams is a property of the prescription. Occupancy is a property of the patient: the proportion of a receptor population actually held by the drug at the moment of measurement. Patients on the same dose of the same agent can sit at very different occupancies, and different agents at their usual doses can sit at the same one. The figures that transfer between drugs are therefore occupancy figures, never milligram equivalences.

The quantity stopped being theoretical when ligand imaging made it measurable. The occupancy figures Kaplan and Sadock supply are striatal and measured by positron emission tomography, and that qualification carries weight twice over. Striatal occupancy is a proxy: the striatum is imaged because it is dopamine-rich and accessible, not because it is the pathway that matters most to the patient. And an occupancy figure is a snapshot of a quantity that rises and falls between doses.

The arithmetic relating drug concentration, receptor affinity and fractional occupancy belongs to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes. What is needed here is developed below and amounts to this: occupancy climbs steeply at first and then flattens, and how firmly an agent grips the receptor decides whether the patient's own dopamine can compete for it.

8.2 One blockade, four consequences, and a pathway that sends no bill

Stahl's Essential Psychopharmacology describes **five** classic dopamine pathways, and this section is built on four consequences rather than on the pathways themselves because the last of them has no established antipsychotic consequence: it arises from several sites, including the periaqueductal grey, projects to the thalamus, and its function is not currently well known.

The **mesolimbic dopamine pathway** runs from the midbrain ventral tegmental area to the nucleus accumbens, territory implicated in pleasurable sensation, in the euphoria of drugs of abuse, and in the delusions and hallucinations of psychosis. The **nigrostriatal dopamine pathway** runs from the substantia nigra to the basal ganglia, forms part of the **extrapyramidal nervous system**, and governs motor function and movement. The **tuberoinfundibular dopamine neurons** run from the hypothalamus to the anterior pituitary, where they are

tonically active in inhibiting prolactin release. The **mesocortical dopamine pathway** leaves the same ventral tegmental area as the mesolimbic projection but ends in prefrontal cortex, where the dorsolateral division may carry cognitive symptoms and the ventromedial division affective ones.

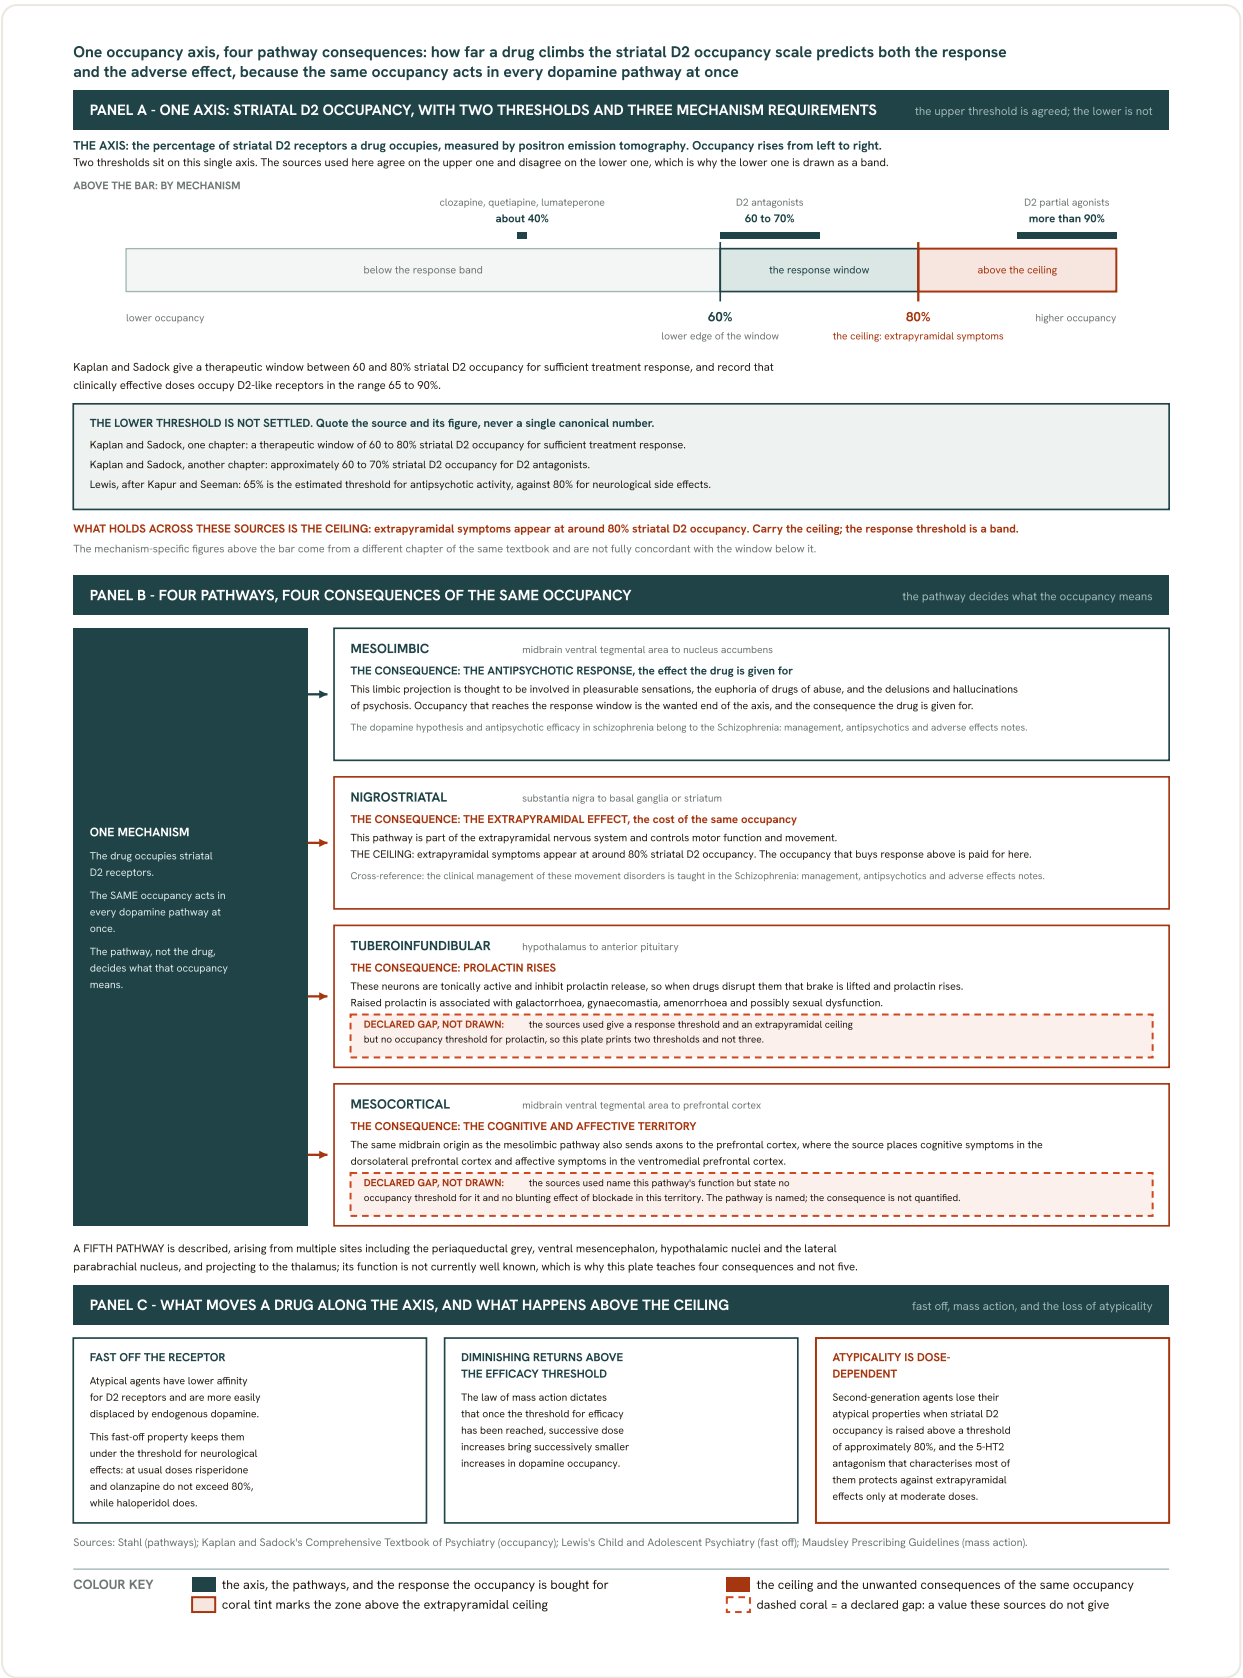


Fig 36.2 **One occupancy axis, four pathway consequences.** How far a drug climbs the striatal D2 occupancy scale predicts both its response and its adverse effects: the sources here agree on a ceiling of around 80% above which extrapyramidal symptoms appear, and disagree on the response threshold, placing it between 60 and 80%. The same occupancy acts at once in the mesolimbic, nigrostriatal, tuberoinfundibular and mesocortical pathways, so the pathway, not the drug, decides what it means. (Stahl; Kaplan and Sadock; Lewis; the Maudsley Prescribing Guidelines.)

Nothing in the pharmacology lets a systemically delivered antagonist prefer one of these destinations to another. The molecule reaches every one of them, and the receptor it blocks is the same receptor in each. This is the single most useful idea in the class, and the accompanying figure carries the same map: the adverse effects are not incidental, they are the same effect occurring somewhere the prescriber did not want it. Everything below follows from spending one occupancy figure in several places at once.

PATHWAY	ORIGIN	PROJECTION	CONSEQUENCE OF D2 BLOCKADE
Mesolimbic	Midbrain ventral tegmental area	Nucleus accumbens	The response the drug is prescribed for
Nigrostriatal	Substantia nigra	Basal ganglia or striatum	The movement-disorder spectrum
Tuberoinfundibular	Hypothalamus	Anterior pituitary	Prolactin rises once tonic inhibition is removed
Mesocortical	Midbrain ventral tegmental area	Prefrontal cortex	Occupancy rises here too, with no established consequence
The remaining pathway	Several sites, including the periaqueductal grey	Thalamus	Not established

MNEMONIC

Benefit, Body, Breast, Brain

The bills a single dopamine blockade pays: **B**enefit where the mesolimbic projection ends, **B**ody movement where the nigrostriatal projection ends, **B**reast effects where the tuberoinfundibular neurons end, and **B**rain, meaning cognition and affect, where the mesocortical projection ends. Recited in that order, the adverse effects arrive attached to their anatomy.

8.3 The mesolimbic consequence: response, and a threshold nobody agrees on

Occupancy at the mesolimbic terminal is what the prescriber is buying, so the natural question is how much of it is enough. The honest answer is that the standard references do not converge, and a candidate carrying one number from one textbook is carrying a number the next textbook contradicts.

Kaplan and Sadock's Comprehensive Textbook of Psychiatry describes a **therapeutic window** lying **between 60 and 80 per cent** striatal D2 receptor occupancy for sufficient treatment response, and reports that clinically effective doses occupy D2-like receptors in the range 65 to 90 per cent. Lewis's Child and Adolescent Psychiatry gives **65 per cent** as the estimated threshold for antipsychotic activity, set against 80 per cent for neurological side effects; that pairing is the 65 to 80 per cent window the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes print. Another chapter of the same comprehensive textbook is more granular still and splits the requirement by mechanism rather than quoting one

figure for the class. A further position in that textbook sets the antagonists' effective occupancy at the level where extrapyramidal effects begin, which read on its own would abolish the window altogether.

Teach it as a band. For an ordinary D2 antagonist the working figure is approximately 60 to 70 per cent; the published range for adequate response runs higher, as far as the point where movement effects start; and none of these positions is settled enough to recite as a single number. What the disagreement means at the bedside is that occupancy predicts response as a band rather than as a switch, so a patient sitting inside the band and not responding is not a patient who needs pushing to the top of it.

8.4 The nigrostriatal consequence: the ceiling everyone agrees on

Same drug, same receptor, different pathway, and now the occupancy is unwanted. Blocking the nigrostriatal projection strips dopaminergic tone from the extrapyramidal system, and past a point the patient's movement changes. That point is the most stable figure on the curve: extrapyramidal symptoms appear above a ceiling of **around 80 per cent** striatal D2 occupancy. It is stable because it recurs unchanged in sources that disagree about nearly everything else here.

Put the response band and the ceiling together and the therapeutic window stops being a phrase and becomes a picture: a stretch of occupancy wide enough to hold a response and narrow enough that a dose increase can walk a patient out of the top of it.

The window also explains something the class labels conceal. Second-generation agents lose their atypical properties once striatal D2 occupancy is raised above a threshold of approximately 80 per cent, and the 5-HT2 antagonism that characterises most of them protects against extrapyramidal effects only at moderate doses. Atypicality is therefore a property of a dose range, not of a molecule. What the typical and atypical labels do and do not predict is taken up elsewhere in these notes, as is the movement-disorder spectrum itself; the clinical management of extrapyramidal effects and tardive dyskinesia belongs to the Schizophrenia: management, antipsychotics and adverse effects notes.

RULE **One occupancy curve produces the response and the movement disorder alike, and the therapeutic window is simply the distance between them.** Response is bought in a band below the extrapyramidal ceiling; extrapyramidal effects begin above it. Every question in this class that sounds like a question about dose is really a question about where on that curve a particular patient sits.

TRAP **Reciting a single number as the occupancy threshold for antipsychotic response.** The published positions run from the bottom of the therapeutic window all the way up to the extrapyramidal ceiling, and one chapter of a standard textbook places the antagonists' effective occupancy at the ceiling itself. Candidates memorise whichever figure their own book printed, which is precisely what a stem offering several plausible percentages is testing. The figure that does not move across sources is the ceiling, and the classic error is to carry it across and offer it as the threshold for response. Give response as a band, name the ceiling separately, and say that the sources differ.

8.5 The tuberoinfundibular consequence: prolactin, and the threshold the sources do not give

Here blockade produces an endocrine result rather than a neurological one, and the mechanism is unusually clean. These neurons are inhibitory and continuously so: their tonic activity holds lactotroph output down rather than switching it off. Remove that brake with a D2 antagonist at the pituitary and prolactin rises. A second standard reference states the same thing from the receptor end, that blockade of pituitary D2 receptors is the basis of prolactin stimulation.

The consequences follow from the hormone rather than from the drug: galactorrhoea, gynaecomastia, amenorrhoea and possibly sexual dysfunction. Notice what that list is made of: most of it is material a patient finds humiliating to volunteer, and the rest is a menstrual change a busy clinic will not think to ask about. The pharmacology predicts the complaint; the consultation has to go looking for it.

What the sources used here do not supply is a number. None of them states a striatal D2 occupancy threshold for hyperprolactinaemia to set beside the response band and the extrapyramidal ceiling, so this section prints two thresholds and no third, and teaches prolactin liability as agent-specific rather than as another point on the shared curve. The Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes do carry a prolactin occupancy figure from the imaging literature; what should not be done is to manufacture one from the figures printed here. Where individual agents fall on prolactin liability is compared elsewhere in these notes.

> **GOOD TO REMEMBER**
>
> Ask about the prolactin consequences by name instead of waiting for them to be volunteered. They follow from the pathway rather than from idiosyncrasy, and breast, menstrual and sexual complaints are exactly the ones patients withhold. A drug-related endocrine effect that is never asked about is commonly answered by silent non-adherence, and the relapse is then read as a failure of efficacy.

8.6 The mesocortical consequence: the one with no readout

This projection ends in the prefrontal divisions named above, and it differs from the other consequences in two ways.

First, it has no bedside measurement. Movement can be examined and prolactin assayed, but no test, threshold or routine scan tells a clinician what occupancy in this pathway is doing to a particular patient. The sources used here give no occupancy figure for it, and none should be quoted from memory.

Second, the projections that most concern the prescriber arise from the same nucleus, which is worth holding in mind: a systemically delivered antagonist cannot raise occupancy at one ventral tegmental terminal without raising it at the other. Whatever is happening in the mesolimbic projection is happening in the mesocortical projection too, and there is no dial that separates them.

The practical consequence is a discipline of attribution: a patient who is flat, slowed and cognitively dulled while taking an antipsychotic has several candidate explanations, and the pharmacology alone cannot separate them from the outside. Naming that as an attribution problem is more honest than asserting a mechanism the sources do not support.

8.7 Affinity and dissociation: why the same occupancy costs different agents different amounts

Occupancy is not something a drug achieves alone; it is a competition with the patient's own dopamine. Agents differ in how tightly they hold the D2 receptor, and that difference has a name and a consequence. Atypical agents have lower affinity for D2 receptors and are more easily displaced by endogenous dopamine, and this **fast-off property** is why, at ordinary doses, they do not sit above the occupancy associated with neurological adverse effects.

The worked example the source gives is exact. 65 per cent D2 occupancy is the estimated threshold for antipsychotic activity, compared with 80 per cent occupancy for neurological side effects, and at usual doses risperidone and olanzapine do not exceed the 80 per cent threshold, whereas haloperidol does. That is the class distinction restated in a currency that can be measured, and the source attributes it to Kapur and Seeman.

Affinity also accounts for the second family of differences, which is by mechanism rather than by label. A plain D2 antagonist needs the band named earlier; a D2 partial agonist needs more than 90 per cent; and clozapine, quetiapine and lumateperone appear to achieve antipsychotic efficacy at approximately 40 per cent postsynaptic striatal D2 occupancy. The occupancy requirement is therefore a property of the mechanism, not of the marketing category. How a partial agonist behaves at the receptor, how the individual agents compare, and clozapine's position outside the ordinary curve are taken up elsewhere in these notes.

TRAP **Assuming one occupancy requirement applies to every agent in the class.** It does not, and the spread is wide enough to invert a stem: a partial agonist needs nearly complete receptor occupancy to work at all, while clozapine and quetiapine work at an occupancy that would be called subtherapeutic for a plain antagonist. The same comprehensive textbook is not fully concordant with itself here either, giving clozapine as 40 to 60 per cent in one place and aripiprazole as about 90 per cent in another. Quote the requirement with its mechanism attached, and never average figures drawn from different chapters.

8.8 Above the threshold: what a dose increase actually buys

The last property of the curve is the one that decides what to do with a partial responder. Occupancy does not rise in proportion to dose. The Maudsley Prescribing Guidelines in Psychiatry puts the reason plainly: the **law of mass action** means the relationship flattens, so that once occupancy has reached the threshold for efficacy, each further increment of dose adds less occupancy than the increment before it.

Read that against the ceiling and the consequence is uncomfortable but clean. Below the threshold, dose increases buy occupancy quickly and buy response with it. Above the threshold, the same increment buys very little further occupancy while carrying the patient towards and then past the point where movement effects begin. Escalating a non-responder is therefore not a

neutral experiment: it moves the patient along the flat part of one curve and up the steep part of another.

Several things follow for practice, and none of them requires a number.

- A patient who has not responded at an adequate dose is usually not a patient who is underdosed, because the occupancy is already there.
- A dose increase made after the threshold has been crossed accepts adverse effects in exchange for very little extra receptor blockade.
- A second-generation agent pushed high enough stops behaving like one, so the property that justified choosing it expires at the dose it is pushed to.
- The occupancy requirement differs by mechanism in a way no dose of a single agent can reproduce, which is a property of the curve rather than a prescribing decision.

PITFALL

Titrating upward in a partial responder until something happens. It is the commonest route by which a well-chosen second-generation agent becomes a poorly tolerated one: the extra milligrams add little occupancy, the patient crosses the extrapyramidal ceiling, and the atypicality that justified the choice is lost at that occupancy anyway. Where to go after an adequate trial has failed belongs to the Schizophrenia: management, antipsychotics and adverse effects notes. Plasma concentrations and the monitoring framework belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

One dopamine blockade lands in four places at once: mesolimbic response, nigrostriatal movement, tuberoinfundibular prolactin, and mesocortical prefrontal cortex with nothing to measure; response is bought in a band the sources place between 60 and 80 per cent striatal D2 occupancy, extrapyramidal effects begin above around 80 per cent, and the sources used here give no third threshold for prolactin.

9 · The movement-disorder spectrum: acute extrapyramidal effects through tardive dyskinesia

Four movement disorders come out of one receptor action, and what separates them is time. Blockade of dopamine receptors in the nigrostriatal pathway generates not a single adverse effect but a sequence, and where a syndrome sits on that sequence predicts more about it than its appearance at the bedside does. A muscular spasm arriving within hours of a first tablet and an involuntary movement disorder declaring itself only after months to years are not unrelated complications to be memorised separately. They are one pharmacological spectrum read at four points on a clock, and the two ends of that clock behave in opposite ways: the early end tracks the dose and remits when the drug is withdrawn, while the late end does neither. Examinations return to this ground because it rewards structure and punishes list-learning: a candidate who has memorised four onset windows without the principle underneath them cannot say why one reaction answers to a dose reduction and another does not. What follows reads the spectrum as

class pharmacology. Which agent to add for an established reaction, and how to manage a patient already disabled by one, belong to the Schizophrenia: management, antipsychotics and adverse effects notes.

About 10 per cent	About 20 per cent	About 25 per cent	5 per cent per year
ACUTE DYSTONIA IN ANTIPSYCHOTIC-TREATED PATIENTS	PSEUDOPARKINSONISM	ACUTE AKATHISIA ON FIRST-GENERATION AGENTS	TARDIVE DYSKINESIA PER YEAR OF EXPOSURE

9.1 Why the clock, and not the movement, is the organising axis

The most useful division in this territory is not between the abnormal movements themselves but between two pharmacological behaviours. **Acute extrapyramidal effects** are usually dose dependent and subside on discontinuing the offending agent. **Chronic extrapyramidal effects** behave in the opposite way: they are not dose dependent, they arise only after prolonged use running to months or years, and they often persist even after the offending medication is stopped. Sort the four syndromes by that division and the topic falls into place; sort them by how the movement looks and you have four descriptions and no principle.

Everything practical follows. If a reaction is dose dependent and reversible, the dose is a genuine lever and time is working for you. If it is neither, the lever was pulled much earlier, at the moment the agent was chosen and the exposure began, which is why the late end is taught as a problem of prevention rather than of treatment. That is not the same as saying dose never mattered, and the two sources here do not read dose the same way. The Indian Psychiatric Society guideline states that chronic extrapyramidal effects are not dose dependent, while the Maudsley lists higher doses among the factors tardive dyskinesia is more common in respect of. Both hold if dose is read at two moments, though neither source states that reconciliation: current dose is not a lever over a syndrome already established, whereas dose across the years of exposure is among the things that raised the chance of acquiring it.

RULE **The acute end of this spectrum is dose dependent and remits on withdrawal; the chronic end is neither.** That one sentence tells you why lowering the dose is a rational answer to a reaction in the first weeks and an insufficient answer to a reaction in the third year. Learn the division before the four syndromes, because the division is what the syndromes are being sorted by.

9.2 The onset windows, and what each one is telling you

Onset is the most discriminating piece of history here, and it is free. **Acute dystonia** is the earliest of the four: it can occur within hours of starting antipsychotic medication, and within minutes if the intramuscular or intravenous route is used. That parenteral compression of the window is not a curiosity: the reaction tracks the arrival of drug at the receptor rather than any slow adaptive process, which is what you would expect of an effect generated by occupancy itself.

The other three windows widen as the underlying process becomes less immediate. Akathisia opens in the same early period but extends well beyond it. **Pseudoparkinsonism** develops days to weeks after the start of antipsychotic medication or an increase in dose, which is slower than a

pure occupancy effect and faster than any structural change. Tardive dyskinesia develops only after months to years, and that long latency is the argument that something other than simple receptor blockade is generating it. Read the windows as a gradient from immediate pharmacology at one end to adaptive change at the other.

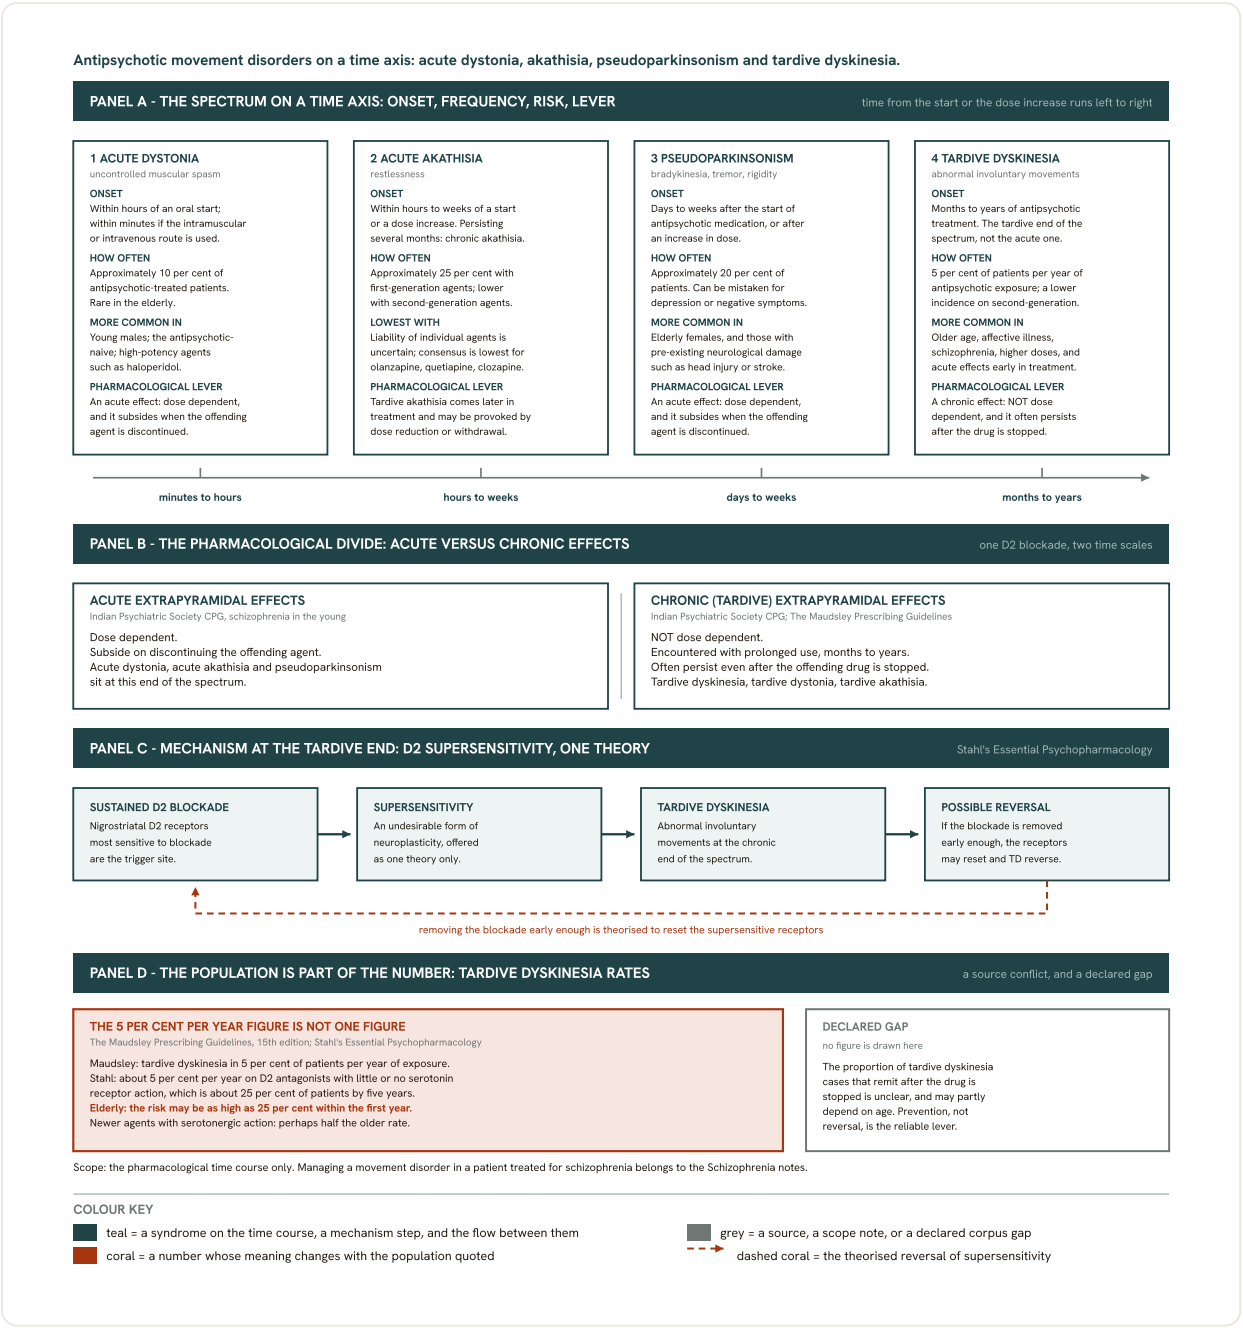


Fig 36.3 **The antipsychotic movement-disorder spectrum read on a time axis.** The four syndromes separate first by when they appear: acute dystonia within hours of an oral start and within minutes by the intramuscular or intravenous route (approximately 10 per cent of patients, more common in young antipsychotic-naive males and with high-potency agents such as haloperidol, rare in the elderly); acute akathisia within hours to weeks of a start or a dose increase (approximately 25 per cent with first-generation agents, lowest with olanzapine, quetiapine and clozapine); pseudoparkinsonism days to weeks after a start or an increase (approximately 20 per cent, more common in elderly females); and tardive dyskinesia after months to years (5 per cent of patients per year of exposure). The acute effects are dose dependent and subside on discontinuing the offending agent, whereas the chronic ones are not dose dependent and often persist after it is stopped, the tardive end being theorised to arise from supersensitivity of the nigrostriatal D2 receptors most sensitive to blockade, which may reset if the blockade is removed early enough. The annual figure carries its population with it: about 25 per cent of patients by five years on older agents, perhaps half that rate on newer serotonergic agents, and as high as 25 per cent in the first year in the elderly, while the proportion of cases that remit after stopping is not stated by the corpus and no figure is drawn for it. (The Maudsley Prescribing Guidelines in Psychiatry, 15th edition; Stahl's Essential Psychopharmacology; Indian Psychiatric Society Clinical Practice Guidelines for the Management of Schizophrenia in Children and Adolescents.)

The figure sets the four syndromes on that axis; the table below adds the two other things a stem supplies, the patient and the agent.

SYNDROME	TIME TO DEVELOP	CLASS BEHAVIOUR	WHERE THE RISK CONCENTRATES
Acute dystonia	Within hours of starting the drug, and within minutes by the intramuscular or intravenous route	Acute: dose dependent, subsides on withdrawal	Young antipsychotic-naive males and high-potency agents; rare in the elderly
Acute akathisia	Hours to weeks after starting the drug or increasing the dose	Acute: dose dependent, subsides on withdrawal	First-generation agents; lowest with olanzapine, quetiapine and clozapine
Pseudoparkinsonism	Days to weeks after the start of treatment or an increase in dose	Acute: dose dependent, subsides on withdrawal	Elderly females and pre-existing neurological damage
Tardive dyskinesia	Months to years of continued treatment	Chronic: not dose dependent, often persists after the drug is stopped	Older age, affective illness and higher doses

MNEMONIC

DAPT: Dystonia, Akathisia, Parkinsonism, Tardive dyskinesia

The four syndromes in the order the clock brings them. **D**ystonia arrives first and arrives fastest by injection, **A**kathisia opens in the same window but runs on for weeks, **P**arkinsonism needs days to weeks to declare itself, and **T**ardive dyskinesia needs months to years. The letters carry onset order and nothing else: nothing about how common each is, and nothing about which disables the patient most.

9.3 Acute dystonia: the reaction that identifies its own risk group

Dystonia is a sustained, uncontrolled muscular spasm: the most dramatic member of the spectrum and the most easily mistaken for something psychiatric when the drug history is not known. The prevalence in the strip above is for antipsychotic-treated patients as a whole, and averaging across a treated population hides how sharply the risk is concentrated. It is more common in young males, in the antipsychotic-naive and with high-potency medications such as haloperidol, while dystonic reactions are rare in the elderly.

Hold that risk profile as a shape rather than a list, because it runs directly opposite to the profile of parkinsonism, which concentrates in older women. A stem specifying the patient's age and sex is usually specifying which of the two it wants. The other half is pharmacological: naivety to the class and high potency both mean a large occupancy step in a system with no opportunity to adapt. Occupancy itself is taught elsewhere in these notes.

9.4 Antipsychotic-induced parkinsonism: the syndrome that hides inside the illness

Pseudoparkinsonism presents with bradykinesia, tremor and muscle rigidity, and the figure in the strip above makes it one of the commonest reasons a patient on maintenance treatment looks unwell without being unwell in the way everyone assumes. It is more common in elderly females

and in those with pre-existing neurological damage such as head injury or stroke. The second half of that sentence is the mechanistically interesting half. A nigrostriatal system already depleted or damaged has less reserve to absorb a given degree of blockade, so the same occupancy produces a clinical picture in one patient and nothing in another. Vulnerability, not dose alone, converts pharmacology into symptoms.

The diagnostic problem this syndrome creates is unique in the spectrum, because its features overlap almost exactly with two things the psychiatrist is independently looking for. The standard prescribing reference warns that pseudoparkinsonism can be mistaken for depression or for the negative symptoms of schizophrenia. All three produce a patient who is slow, flat and hard to engage. The discriminator is not the mental state examination but the timing and the motor signs.

GOOD TO REMEMBER

Before adding an antidepressant to a patient who has become slow and withdrawn on an antipsychotic, account for the drug already prescribed. Establish when the change began relative to the last dose adjustment, and examine for tone and tremor, because a picture that began days to weeks after a dose increase is a drug effect until shown otherwise. Treating it as a mood problem adds a second drug to a problem the first one caused.

9.5 Akathisia: one word covering three different phenomena

Akathisia is the member of the spectrum most often missed, and the only one whose defining feature is subjective. It is an inner restlessness the patient localises to the body and relieves by moving, and it is common: approximately 25 per cent of patients on first-generation antipsychotics and less often on second-generation agents. The source warns explicitly that studies vary widely, so this is a magnitude rather than a precise rate. Unusually for this spectrum, the word covers three phenomena separated by time.

- **Acute akathisia** occurs within hours to weeks of starting an antipsychotic or increasing the dose, which places it alongside dystonia at the early end of the clock.
- **Chronic akathisia** is the name given to akathisia that has persisted for several months, and it is a duration label rather than a separate mechanism.
- **Tardive akathisia** tends to occur later in treatment, and may be provoked by dose reduction or withdrawal of the antipsychotic.

That last point is the one worth carrying, because it inverts the reflex the acute end of the spectrum trains. At the acute end, restlessness on a rising dose means the dose is too high; at the tardive end, restlessness may appear precisely because the dose is being lowered or the drug stopped. Withdrawal-emergent movement is a general property of the chronic end of this spectrum, which is why a taper is a period of active observation, not one in which adverse effects can be assumed to be receding.

PITFALL

Reading akathisia as psychotic agitation and answering it by raising the dose. The source giving the incidence figure warns that akathisia can be mistaken for psychotic agitation, and the error is self-reinforcing: the drug that caused the restlessness is increased, the restlessness worsens, and the worsening is taken as proof that the dose was inadequate. The discriminator is subjective and has to be asked for: agitation is driven by the mental state, whereas this restlessness is bodily. Ask before the dose is changed, not after.

9.6 Tardive dyskinesia: the number, and the population it belongs to

Tardive dyskinesia is the abnormal involuntary movement disorder at the chronic end of the spectrum, and it is quantified as a rate rather than a prevalence because it accrues with exposure. It occurs in **5 per cent of patients per year of antipsychotic exposure**. The unit is the whole point: this is per cent of patients per year of exposure, not a lifetime figure, so it compounds in a way a prevalence does not. Risk is stated as higher in relation to several factors.

- Older age, and higher doses. This is the entry the dose-independence statement has to be read against.
- Schizophrenia, and also affective illness, which is a reminder that this risk is not confined to patients with psychotic disorders.
- Acute extrapyramidal effects early in treatment, clinically the most useful of the set, because it means the early end of the spectrum carries prognostic information about the late end.

The rate itself is the most misquoted number here, because two standard sources attach the same figure to different populations.

STATED BY	THE POPULATION THE FIGURE APPLIES TO	THE FIGURE
The Maudsley Prescribing Guidelines in Psychiatry	Antipsychotic-treated patients as a whole, with a lower incidence in those on second-generation agents	5 per cent of patients per year of antipsychotic exposure
Stahl's Essential Psychopharmacology	Patients maintained on D2 antagonists with little or no serotonin receptor action	About 5 per cent every year, and about half that rate for the newer agents with serotonergic action

Stahl carries the arithmetic forward, and the forward projection is what makes the figure alarming rather than academic: about **25 per cent of patients by five years**. He sets that against an illness beginning in the early twenties and requiring lifelong treatment, which is why a rate that sounds tolerable in a single year is not a tolerable prospect over a treatment career. Age changes the picture sharply again: the risk in elderly subjects may be **as high as 25 per cent within the first year of exposure** to those same agents. An older patient therefore reaches in one year what a general population reaches in five.

TRAP **Quoting the annual tardive dyskinesia rate without saying whose rate it is.** The figure is safe as a headline for the older agents and it is not safe as a universal figure for every antipsychotic, because one source attaches it to antipsychotic-treated patients as a whole while the other restricts it to D2 antagonists with little or no serotonin receptor action and halves it for the newer serotonergic agents. It is badly wrong for the elderly, in whom the first-year risk is stated as up to a quarter. Give the rate with the population it was measured in, or give the direction and decline the number.

9.7 Supersensitivity: why the late end behaves unlike the early end

A mechanism has to explain three things that simple receptor blockade does not: the long latency, the absence of dose dependence, and persistence after the drug has gone. One theory holds that the nigrostriatal D2 receptors most sensitive to blockade trigger an undesirable form of neuroplasticity called **supersensitivity** in reaction to that blockade. The word theory is the source's own and should be kept: this is a proposal that accounts well for the phenomenon, not an established mechanism.

Read that way, the pieces cohere. An adaptive change in receptor number or sensitivity takes time to develop, which gives the latency; it is driven by sustained rather than instantaneous blockade, which is why current dose predicts it poorly; and an adaptation already made does not disappear when the stimulus does, which is why the syndrome outlasts the drug. The same logic predicts withdrawal-emergent movement: remove the antagonist from a supersensitive system and the movements may appear or worsen before they settle. The receptor-regulation science underneath this belongs to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

The theory carries one hopeful corollary, and it is the reason early recognition matters. If the blockade is removed early enough, tardive dyskinesia may reverse, through a resetting of those supersensitive receptors. How often that happens is exactly where the evidence stops. The Maudsley Prescribing Guidelines in Psychiatry states plainly that the proportion of cases showing reversibility on cessation of antipsychotic medication is unclear and may partly depend on age. That blank is a stronger teaching point than any figure would be: nobody can promise a patient that stopping the drug will undo what it has already done.

9.8 Within-class liability: what actually moves the risk

If prevention is the lever, the question becomes which agent, and here the honest answer has two halves. The first half is that the class-level direction is consistent: the incidence of tardive dyskinesia is lower in those on second-generation agents, and for agents with serotonergic receptor action the rate may be about half the rate of the older drugs. Serotonergic action is the axis those comparisons turn on, which is why the generation label is a proxy for a receptor property rather than a property in itself.

The second half is that agent-by-agent precision is not available. For akathisia the sources say directly that the relative liability of individual antipsychotic medications is uncertain, while still recording a consensus that the incidence is lowest for olanzapine, quetiapine and clozapine. Hold those two statements together: there is a defensible bottom to the ranking and no defensible

middle. Ranking the class across the other adverse-effect axes is taken up elsewhere in these notes, and the drug emergencies adjacent to this territory, including neuroleptic malignant syndrome, belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

IN INDIA

The acute versus chronic division that organises this topic is set out for Indian practice by the Indian Psychiatric Society, in its clinical practice guideline for the management of schizophrenia in children and adolescents; the source population there is paediatric, but the pharmacological division is class level. While the picture is acute, the dose is a real lever, because acute effects are dose dependent and subside on discontinuing the offending agent. Once it is tardive that lever is gone, so the decision that mattered was the one taken at initiation: in a young antipsychotic-naive patient a high-potency agent carries the dystonia risk, and for akathisia the recorded consensus puts olanzapine, quetiapine and clozapine lowest. What to do about an established reaction belongs to the Schizophrenia: management, antipsychotics and adverse effects notes.

Acute extrapyramidal effects are dose dependent and subside when the drug is withdrawn; the tardive ones are dose independent, take months to years to appear, often persist after withdrawal, and their annual incidence figure belongs to the population it was measured in, not to every antipsychotic.

10 · Clozapine: what makes it different and what that costs

Every other agent in this class can be understood from a single occupancy curve, and clozapine cannot. It is effective at a striatal dopamine occupancy the rest of the class would call subtherapeutic, it carries an adverse-effect burden no other psychotropic in routine use matches, and it is the only drug in psychiatry whose continued supply is conditional on a blood result. Those three statements are one: a receptor profile unusually broad and unusually weak at the receptor the class is built around. What follows reads clozapine as a pharmacological object rather than as a treatment. Which patients should receive it, how treatment resistance is defined, and what to add when it works only partly belong to the Schizophrenia: management, antipsychotics and adverse effects notes. Read this way, the monitoring apparatus stops being an arbitrary list and becomes reconstructible.

0.9 per cent	Up to 3.0 per cent	32 to 60 per cent	15.0 to 27.5 per cent
META-ANALYTIC INCIDENCE OF SEVERE NEUTROPENIA	MYOCARDITIS, EARLY AND EASILY MISSED	REPORTED CONSTIPATION	FATALITY ONCE ILEUS OCCURS

10.1 The occupancy anomaly that defines the drug

The rule for the class is that a dopamine receptor antagonist becomes effective only when it occupies most of its striatal target, and the sources disagree about how much. The chapter grounding this section takes the highest position, placing the requirement at approximately 80 per cent of striatal D2 receptors for the antagonists as a group on positron emission tomography

with highly selective ligands, with occupancy above that buying parkinsonian adverse effects rather than extra benefit. That figure is not the settled response threshold. Other chapters put the antagonist requirement lower, and the number that stays still across sources is the extrapyramidal ceiling rather than the point response begins; the occupancy section of these notes teaches response as a band. Clozapine undercuts every version of the requirement: it is effective when only **40 to 60 per cent** of those same receptors are occupied, and the same textbook gives a lower figure still for it elsewhere. The mechanism that accounts for the rest of the class does not account for this one, which is why clozapine cannot be taught as one more entry in an alphabetical list of atypicals.

Why it works down there is explicitly unsettled; two explanations are proposed. The first is that activity at receptors other than the dopamine target, particularly 5-HT2A, contributes to the antipsychotic effect in its own right. The second involves the major active metabolite, **N-desmethylclozapine**, a partial agonist at D2 and D3 receptors and a strong inverse agonist at 5-HT2A and 5-HT2C. Since a partial agonist behaves differently from an antagonist at the same receptor, parent compound and metabolite are not pharmacologically interchangeable. The adverse effects generated by high striatal blockade are therefore not the ones limiting this drug.

RULE **Clozapine is the antipsychotic that works below the occupancy the rest of the class requires.** Efficacy is therefore not being bought with dopamine blockade; it is bought elsewhere in the receptor profile, and every other receptor the molecule touches is a potential adverse effect. That sentence predicts this drug's burden before any adverse effect is listed.

10.2 A promiscuous profile, and what breadth costs

Selectivity is what makes a drug predictable, and clozapine is not selective: it engages dopaminergic, serotonergic, muscarinic, histaminergic and adrenergic targets, and a patient's clinical picture is largely the sum of those engagements. Reading a receptor profile forward into an adverse-effect signature is taught elsewhere in these notes.

The most quantitatively striking of these actions is at the muscarinic receptor. Clozapine is **strongly antimuscarinic**, to the extent that 50 to 100 mg of oral clozapine equals the anticholinergic potency of 1 mg of benztropine, and constipation on this drug is therefore predictable pharmacology rather than an incidental complaint. Two cautions attach. The equivalence concerns anticholinergic potency only and says nothing about antipsychotic dose equivalence, and adding any strongly anticholinergic agent is not a neutral act: the risk of ileus is doubled by concurrent use of strongly anticholinergic agents. Interactions across drug classes belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

MNEMONIC

Marrow, Muscle, Mind, Motility

The organ systems carrying clozapine's serious burden: **M**arrow (neutropenia), cardiac **M**uscle (myocarditis), **M**ind (seizures) and gut **M**otility (constipation progressing to ileus). The one monitored obsessively is not the one with the highest case fatality.

10.3 Exposure: the enzyme and the ratio that reads it

Clozapine is metabolised by several cytochrome isoenzymes, but not equally. CYP1A2 is the most important isoenzyme in clozapine metabolism, and especially so at lower plasma concentrations. That qualification matters: the enzyme dominating at the bottom of the range is the one whose induction or inhibition moves a patient furthest and fastest. The enzymology itself belongs to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

What is specific to this drug is that its metabolism can be read off a paired assay. The ratio of clozapine to norclozapine, the **metabolic ratio**, reads **1.32** in a patient who is an extensive metaboliser at the relevant isoenzymes and exposed to neither inhibitor nor inducer, a reference value derived from nearly 5000 samples obtained from over 2000 patients. It is dimensionless, so it escapes the unit confusion afflicting absolute levels, and reads in three bands.

- A ratio near the reference value: metabolism is behaving as expected and the sample was taken at trough.
- A ratio significantly above the reference value, for example above 2.00: a non-trough sample, slow metabolism, or an inhibitor at work.
- A ratio at or below 1.00: induction.

Induction is the direction that bites most often, because tobacco smoke induces this enzyme and many patients on clozapine smoke. The dangerous moment is the stopping, by choice or on admission to a smoke-free ward, since induction is withdrawn and exposure climbs. The standard references disagree twice here, and the disagreement is itself the teaching point: on how few cigarettes a day fully induce the enzyme, and on how large the rise in level is on cessation and therefore how big the first dose reduction should be. The cytochrome section of these notes quotes one reference for both; the other puts the cigarette count higher and the dose step differently, so either figure travels with its book. They agree on direction and on method: induction is maximal well below heavy smoking, and the level should be measured across the transition rather than predicted. Target plasma concentrations and the monitoring framework belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

PITFALL

Reading a high metabolic ratio as proof of slow metabolism. A ratio above 2.00 is equally well produced by a sample drawn at the wrong time, because parent drug and metabolite do not fall in parallel after a dose. Establish that the sample was a genuine trough before concluding anything about enzymes or reaching for an interaction.

10.4 The haematological hazard: its true size and its shape

The event that built the entire monitoring apparatus was classically called **agranulocytosis**. In the United States that term has been replaced for regulatory purposes by **severe neutropenia**, defined by a count below 500/mm3. The renaming is not cosmetic: it moves the definition onto a single measurable variable and away from a clinical impression, which is why modern schemes speak in counts rather than syndrome names.

Its size has been estimated many times and the estimates do not sit still. A meta-analysis of 108 publications published in 2018 gives the figure quoted above, with a 95 per cent confidence interval of 0.7 to 1.1 per cent, against a spread across individual studies of 0.38 to 2.0 per cent. That spread reflects different schemes and different eras rather than different biology. Death is rarer by a further order of magnitude: with modern monitoring and treatment the incidence of death related to neutropenia during clozapine use is estimated at **0.013 per cent**, with a 95 per cent confidence interval of 0.01 to 0.017 per cent, roughly seventy-fold smaller than the incidence of severe neutropenia itself. The gap between event and death is the monitoring working.

The hazard also has a shape, and the schedule is built on it. Risk is front-loaded: peak incidence of severe neutropenia falls at 1 month of exposure and the hazard becomes negligible after 1 year of treatment, with the period of greatest risk occupying the first six months. Recovery once the drug is withdrawn is usually brisk: the **median time to resolution of severe neutropenia is 12 days**. Read that as a median, not a maximum, and read the curve as an argument for front-loaded testing rather than permission to stop.

BAND	ABSOLUTE NEUTROPHIL COUNT	WHAT IT TRIGGERS
Mild neutropenia	1000 to 1499/mm3	Continue treatment; count three times weekly until at or above 1500/mm3, then return to the previous interval
Moderate neutropenia	500 to 999/mm3	Interrupt treatment, with haematology consultation recommended; daily counts until at or above 1000/mm3, then three times weekly until at or above 1500/mm3, then weekly for a further four weeks
Severe neutropenia	below 500/mm3	Interrupt treatment, with haematology consultation recommended; do not rechallenge unless the prescriber judges the benefits to outweigh the risks

10.5 The monitoring apparatus, and why two countries built it differently

Two national schemes govern this drug, they do not agree, and the disagreement is instructive rather than embarrassing. The United Kingdom tracks two variables, the total white cell count and the neutrophil count, banding the result green, amber or red. The United States tracks the **absolute neutrophil count** alone. Neither is wrong; they differ on how much redundancy is worth its false alarms. Both make explicit allowance for **benign ethnic neutropenia**, a constitutionally low neutrophil count in an otherwise well person, by setting a lower entry threshold rather than by excluding those patients.

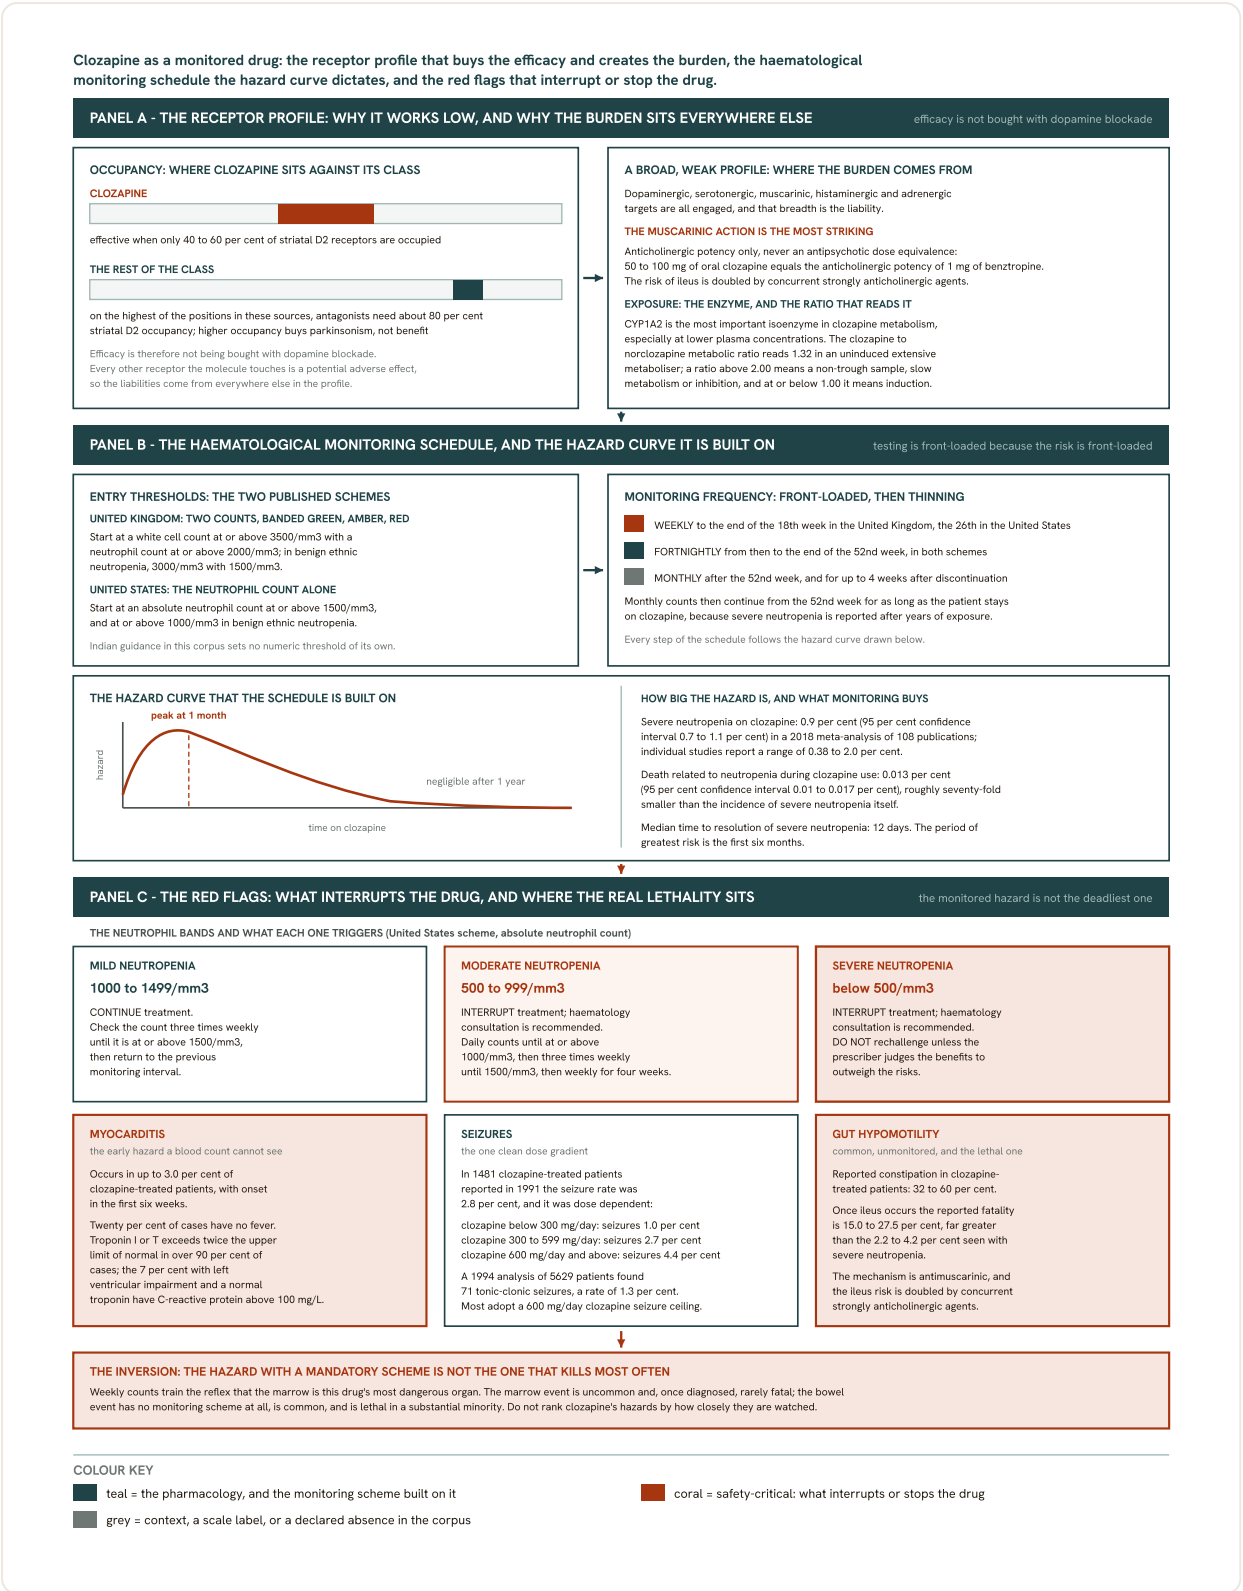


Fig 36.4 Clozapine: the profile that buys the efficacy, the monitoring schedule the hazard curve dictates, and the red flags that stop the drug. Clozapine is effective when only 40 to 60 per cent of striatal D2 receptors are occupied, where the antagonists as a group need approximately 80 per cent, so the efficacy is bought elsewhere in a broad receptor profile and the burden is paid in marrow, heart, brain and bowel. The neutrophil bands and the weekly, then fortnightly, then monthly schedule track a hazard that peaks at 1 month and becomes negligible after 1 year, yet monthly counts never stop; myocarditis, seizures and gut hypomotility sit entirely outside what a blood count can see. Indian guidance in this corpus prints no numeric threshold of its own, so one published scheme is adopted in full and recorded. (Stahl's The Clozapine Handbook; Kaplan and Sadock's Comprehensive Textbook of Psychiatry 2024.)

ELEMENT OF THE SCHEME	UNITED KINGDOM	UNITED STATES
Counts tracked	White cell count and neutrophil count together, banded green, amber and red	The absolute neutrophil count alone
Entry threshold, general population	White cell count at or above 3500/mm3 with a neutrophil count at or above 2000/mm3	Neutrophil count at or above 1500/mm3
Entry threshold in benign ethnic neutropenia	White cell count at or above 3000/mm3 with a neutrophil count at or above 1500/mm3	Neutrophil count at or above 1000/mm3
End of the weekly phase	Runs to the end of the 18th week	Runs to the end of the 26th week
Fortnightly and monthly phases	Fortnightly to the end of the 52nd week, monthly thereafter	Identical to the scheme alongside from this point onward

Every element of that table derives from the hazard curve rather than from regulation for its own sake: testing is weekly while risk peaks, thins as the hazard falls, and settles at monthly once the curve flattens. Both schemes also continue counts for up to four weeks after the drug is discontinued, because a marrow effect already initiated does not stop when the tablets do. The element that does not derive from the curve is why monitoring never ends: monthly testing continues from the 52nd week onwards for as long as the patient remains on clozapine, because severe neutropenia has been reported after years of exposure. A negligible hazard is not a zero hazard.

IN INDIA

Neither published scheme is Indian, and Indian guidance directs that the haemogram be monitored according to established guidelines without setting a numeric threshold of its own. Two decisions follow at the point of prescribing. Adopt one published scheme in full and document which, because the two differ in what is counted and in when the weekly phase ends, and a service that blends them produces counts nobody can interpret. Then take a baseline count and consider benign ethnic neutropenia before treating a low result as a bar to treatment, since both schemes set a lower entry threshold for that situation rather than withholding the drug.

10.6 Myocarditis: the early hazard that does not announce itself

Myocarditis defeats a monitoring scheme built around the marrow, because it is early, invisible to a neutrophil count and unhelpfully bland in presentation. **Myocarditis** on clozapine has an onset during the first six weeks, with rare exceptions, which places it inside the window when everyone is watching the blood and nobody the heart. The picture is malaise or an influenza-like illness, often with chest pain, and the expectation of fever is the trap: twenty per cent of cases have no fever at all.

The laboratory signature is more dependable, and it has two components because either alone misses a subgroup. The troponin I or T level exceeds twice the upper limit of normal in over 90 per cent of cases, which makes troponin the single most useful test. But 7 per cent of cases with

left ventricular impairment on echocardiography have a normal troponin together with a C-reactive protein above 100 mg/L, so a normal troponin in a symptomatic patient in the early weeks does not close the question. Recognition and acute management of this and the other psychotropic emergencies belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

GOOD TO REMEMBER

In a patient in the first weeks of clozapine treatment reporting malaise, an influenza-like illness or chest pain, send troponin and C-reactive protein together and do not wait for a temperature, since fever is absent in a substantial minority of cases. The weekly blood count will not detect this complication.

10.7 Seizures: the one adverse effect with a clean dose gradient

Most of clozapine's adverse effects are receptor-mediated and present across the exposure range. Seizures behave differently and are its clearest concentration-dependent liability. The first large data set, covering 1481 clozapine-treated patients and published in 1991, reported an overall seizure rate of 2.8 per cent, and within it the gradient was orderly: 1.0 per cent below 300 mg/day, 2.7 per cent at 300 to 599 mg/day, and 4.4 per cent at 600 mg/day and above. The same authors estimated cumulative seizure risk in that group to be as high as 10 per cent at 3.8 years, a reminder that a cross-sectional rate and a cumulative risk are different quantities.

The headline rate does not replicate. A subsequent analysis by the same authors, of 5629 patients in 1994, found only 71 tonic-clonic seizures, a rate of 1.3 per cent, and later case series with variable exposure periods report figures higher than either. The sources state plainly that no dose or plasma-level threshold for clozapine-associated seizures is well defined, and that in the absence of one most practitioners have adopted a theoretical ceiling of 600 mg/day for seizure risk. That ceiling is a convention derived from the gradient, not a validated threshold.

TRAP **Quoting a single seizure rate for clozapine.** The same investigators reported 2.8 per cent in one cohort and 1.3 per cent in a larger one three years later, and neither answers a bare question about incidence. Candidates memorise whichever figure their own textbook printed. The reproducible finding is not the headline rate but the dose gradient, which is what a stem describing a patient at higher exposure is actually testing. Give the rate with its cohort and its denominator, or give the gradient.

10.8 Gastrointestinal hypomotility: the common hazard and the lethal one

Constipation on clozapine is not a nuisance effect; it is the visible end of **gastrointestinal hypomotility**, an antimuscarinic action on gut motility that can progress to ileus. Its frequency is on a different scale from the haematological event, and the reported incidence quoted above varies that widely, the source explains, because of the variety of clinical criteria used across studies rather than because populations differ.

Frequency alone would not make it important; lethality does. Once ileus occurs its reported fatality is **15.0 to 27.5 per cent**, far greater than the 2.2 to 4.2 per cent seen with severe

neutropenia. Put the two together and the risk profile inverts: the hazard with a mandatory testing scheme is uncommon and rarely fatal, while the one with no scheme at all is common and, once declared, lethal in a substantial minority. The mechanism being antimuscarinic, the levers are the ones that made constipation predictable: anticipate it, ask about it actively, and avoid stacking further anticholinergic burden.

TRAP **Ranking clozapine's hazards by how closely they are monitored.** Weekly counts train the reflex that the marrow is this drug's most dangerous organ; the case fatality figures point the other way. Note too that the standard reference does not agree with itself here: where the haematology is discussed it gives a case fatality of 2.1 per cent, with a 95 per cent confidence interval of 1.6 to 2.8 per cent, and where the bowel is discussed it gives the wider range quoted above. Do not average them, and do not present either as the risk of dying on clozapine, which is a different and far smaller quantity: case fatality is conditional on the event having already happened.

Clozapine works at 40 to 60 per cent striatal D2 occupancy, below the band the rest of the class needs, and pays for that breadth in marrow, heart, brain and bowel, of which only the marrow is monitored by mandate and only the bowel kills in double figures.

11 · Long-acting injectables: the pharmacokinetics of switching

A long-acting injectable is an ordinary antipsychotic molecule with its absorption deliberately made the slowest step in its own kinetics, and nearly every error made with this group follows from forgetting that. Nothing is added to the pharmacodynamics. The receptor profile, the occupancy curve and the adverse-effect signature remain those of the parent drug. What changes is the front of the concentration curve: a depot surrenders its drug over weeks rather than over hours, so the rate of release from the injection site, not the rate of elimination, governs everything the prescriber sees. Every practical question that follows is therefore a question about release.

The decision to offer a long-acting injectable at all, and what it buys in adherence and relapse prevention, belongs to the Schizophrenia: management, antipsychotics and adverse effects notes. The general behaviour of half-life and steady state is developed elsewhere in these notes, and the kinetic equations underneath it belong to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes. What is taught here is the arithmetic a depot makes visible and an oral drug conceals: switching, loading, covering and waiting.

3	6 to 8 weeks	about 28 days	about 52 weeks
HALF-LIVES TO NEARLY ALL OF STEADY STATE	MINIMUM BEFORE A DEPOT IS AT STEADY STATE	TO PEAK FOR RISPERIDONE MICROSPHERES	TO STEADY STATE FOR THREE-MONTHLY PALIPERIDONE

11.1 The vehicle decides the kinetics, and it decides the test dose

Three release architectures cover almost everything in current use, and they are not interchangeable:

- An ester of the parent drug dissolved in **an oil vehicle**, which is what the first-generation decanoates are. The ester must be liberated from the oil and hydrolysed back to the parent compound before anything happens.
- Parent drug coated in polymer to form **microspheres** suspended in water, which is what risperidone long-acting injection is. There is no ester at all.
- An **ester prodrug formulated as an aqueous nanosuspension**, which is what paliperidone palmitate is. The chemistry of the first architecture is combined with the vehicle of the second.

The vehicle also settles a question candidates are routinely asked in the wrong terms. For a first-generation injectable the **test dose** is a small quantity of active drug in a small volume of oil, and it serves two purposes at once: it tests the patient's sensitivity to extrapyramidal effects, and any sensitivity to the base oil itself. For second-generation injectables a test dose may not be required, and the reason is a double one, because extrapyramidal propensity is lower and the aqueous base is not known to be allergenic. The test dose is therefore a property of the preparation, not a ritual owed to every injection.

11.2 Accumulation, and why the opening weeks cannot be judged

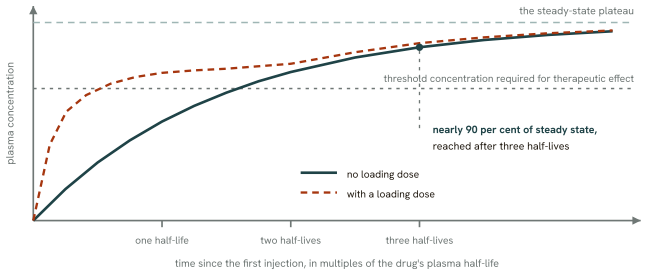
Nothing about a depot's opening months behaves the way a prescriber's instinct expects. With most long-acting injectable preparations, plasma concentrations climb for several weeks to months with no increase in dose at all, because each injection is given before the previous one has been cleared. That is **accumulation**, and its consequence is that the patient reviewed at two months is on a materially different exposure from the patient reviewed in the first week, on an unchanged prescription.

The arithmetic underneath is fixed and not negotiable. Attainment of steady state follows logarithmic rather than linear characteristics, so **nearly 90 per cent** of steady-state concentrations are present after three half-lives, and the time taken to arrive there is independent of both the dose and the dosing frequency. A **loading dose** is the only real exception, and it is a narrower one than it appears: it produces prompt therapeutic plasma concentrations, but the time to steady state remains exactly what it was. Note also that steady state and the concentration required for a therapeutic effect are two different things, and a preparation can be doing useful work long before it plateaus.

Long-acting injectables: how steady state accrues, what a loading dose can and cannot do, and why the interval is what it is

PANEL A - HOW A DEPOT FILLS UP: STEADY STATE ACCRUES LOGARITHMICALLY

a loading dose does not hurry it



THE ARITHMETIC OF STEADY STATE

Maudsley Prescribing Guidelines, 15th edition

Accrual is logarithmic, not linear: nearly 90 per cent of steady state is reached in three half-lives.

Time to steady state is independent of dose and of dosing frequency. You cannot hurry it up by giving more, more often.

A loading dose produces prompt therapeutic plasma levels, but the time to steady state stays the same.

At steady state, troughs just before and after a dose sit substantially above the threshold concentration needed for therapeutic effect.

PANEL B - WHAT THE CLOCK MEANS AT THE BEDSIDE

do not escalate early, do not shorten the interval

THE ACCUMULATION CLOCK

Maudsley, 14th edition

Plasma levels rise over several weeks to months with NO increase in dose. This is accumulation.

Steady state is only achieved after at least 6 to 8 weeks.

An early rise is the drug piling up, not the dose being right.

DO NOT ESCALATE EARLY

Maudsley, 14th edition

Increase a dose only after careful assessment over at least one month, and preferably longer.

Dose increases inside this window are illogical and cannot be judged.

Doses may be reduced if adverse effects occur.

DO NOT SHORTEN THE INTERVAL

Maudsley, 14th edition

Every long-acting injectable can be given safely at its licensed interval. There is no evidence that shorter intervals improve efficacy.

Levels keep falling for hours or days AFTER a dose, so a jump in wellbeing right after one makes no kinetic sense.

THE TEST DOSE

Maudsley, 14th edition

First-generation, in oil: a test dose of active drug in a small volume of oil tests sensitivity to extrapyramidal effects AND to the base oil.

Second-generation, aqueous: often not needed, since extrapyramidal risk is lower and the base is not allergenic.

PANEL C - THE FORMULATIONS COMPARED ON THE SAME THREE MEASURES

time to peak, half-life, time to steady state

FORMULATION	TIME TO PEAK	PLASMA HALF-LIFE	STEADY STATE	WHAT IT MEANS IN THE CLINIC
Fluphenazine decanoate (oil)	8 to 12 days	7 to 10 days	about 8 weeks	Tabulated peak is disputed; two peaks likely.
Haloperidol decanoate (oil)	3 to 9 days	21 days	about 14 weeks	Peaks within days, settles only after months.
Risperidone microspheres	about 28 days	3 to 6 days	about 8 weeks	Fixed 2-weekly interval, no longer allowed.
Olanzapine pamoate	2 to 6 days	30 days	about 12 weeks	Carries the acute hazard in Panel D.
Aripiprazole one-monthly	4 days deltoid; 5 to 7 gluteal	27 days	about 20 weeks	Judge the response late, not early.
Paliperidone palmitate, monthly	13 days	25 to 49 days	about 20 weeks	Loaded on days 1 and 8, so no oral cover.
Paliperidone palmitate, 3-monthly	recorded inconsistently in the source	84 to 95 days deltoid; 118 to 139 days gluteal	about 52 weeks	The longest time to steady state here. Switch only from a stable patient.
Paliperidone palmitate, 6-monthly	29 days at the lower dose; 32 days at the higher dose	148 days at the lower dose; 159 days at the higher dose	not known	The kinetics now run in months. Gluteal route only.

PANEL D - LOADING, THE INTERVAL LADDER, AND THE ONE ACUTE HAZARD

every dose is named with its indication

THE LOADING REQUIREMENT: PALIPERIDONE PALMITATE LOADS, RISPERIDONE MICROSPHERES CANNOT

Paliperidone palmitate, antipsychotic maintenance, day 1: 150 mg intramuscularly into the deltoid; day 8 (plus or minus 4 days): 100 mg into the deltoid.
Then, for antipsychotic maintenance, 50 to 150 mg monthly (plus or minus 7 days), into the deltoid or the gluteal muscle.
Both loading injections are deltoid because a single deltoid dose gives on average a 28 per cent higher peak concentration than the same gluteal dose.
Risperidone microspheres cannot be loaded: 3 to 4 weeks pass before the first injection produces therapeutic plasma levels, so the previous antipsychotic is continued at full dose for at least 3 weeks, and sometimes for 6 to 8 weeks; a test dose is neither required nor sensible.

STRETCHING THE INTERVAL, AND THE PRICE OF A MISSED INJECTION

To the three-monthly formulation: 4 months or more on the one-monthly with the last two doses the same, then give it in place of the next scheduled dose.
For antipsychotic maintenance the three-monthly dose is 3.5 times the monthly, not 3 times: 50 mg becomes 175 mg, 75 mg becomes 263 mg, 100 mg becomes 350 mg, 150 mg becomes 525 mg.
For antipsychotic maintenance the six-monthly needs 100 mg or 150 mg monthly for 4 months or more, or one 350 mg or 525 mg three-monthly dose, and is gluteal only.
The deltoid raises the peak by 11 to 12 per cent for the three-monthly formulation, against 28 per cent for the one-monthly.
Correct dosing is critical: patients receiving fewer than 12 injections a year have an increased risk of relapse.
A second reading in the same source places the cut at 10 doses a year or fewer, against 11 or 12 doses.

THE ONE ACUTE HAZARD: POST-INJECTION DELIRIUM SEDATION SYNDROME WITH OLANZAPINE PAMOATE

Thought to occur when the pamoate salt meets a large volume of blood or plasma and releases a large amount of olanzapine at once, so plasma levels may exceed 800 microg/L.
Incidence is under 0.1 per cent of injections, and 0.044 per cent, under 1 in 2,000, in one study; 86 per cent of reactions occur within 1 hour, mean 30 minutes.
In that study 91 per cent appeared within 1 hour; reactions resolve fully within 72 hours, and at least 3 hours of observation is mandated in the European Union and United Kingdom.

COLOUR KEY

- teal = the kinetics, the formulations and the arithmetic
- tinted = alternate table rows and the summary bands
- coral = the hazard, the caution and the number to get right
- grey = axes, sources and scope notes

Fig 36.5 Long-acting injectable release kinetics. Steady state accrues logarithmically: nearly 90 per cent of it is reached in three half-lives, and the time to steady state is independent of dose and of dosing frequency, so a loading dose buys prompt therapeutic plasma levels without shortening the climb. Levels rise for weeks to months with no dose increase, steady state needs at least six to eight weeks, and the dose should rise only after assessment over at least a month. Paliperidone palmitate loads on days 1 and 8 into the deltoid, which peaks higher than the gluteal route; risperidone microspheres cannot be loaded. (The Maudsley Prescribing Guidelines in Psychiatry, 14th and 15th editions.)

For a depot those two facts set a floor. Steady state is only achieved after **at least 6 to 8 weeks**, and a dose should be increased only after careful assessment over at least one month, and preferably longer. A dose increase inside that window is not merely premature; it is uninterpretable, because the concentration was going to rise anyway and no clinician can afterwards say which of the two changes produced the effect.

RULE **Time to steady state is fixed by the preparation's own release and elimination kinetics; neither the size of the dose nor the frequency of dosing can move it.** That single sentence disposes of most depot questions. It explains why the licensed interval is what it is, why the first month of treatment is uninterpretable, why a loading dose buys height and not time, and why the correct response to a partial response early in treatment is to wait rather than to escalate.

Three half-lives take a long-acting injectable to nearly all of its steady state, and a loading dose changes how high the early concentrations run, never when steady state arrives.

TRAP **Treating a loading dose, or a shortened interval, as a way of reaching steady state sooner.** The error is a conflation of two distinct quantities: a prompt therapeutic concentration, which a loading dose genuinely delivers, and steady state, which it does not touch. Stems offer a patient only partly better after a few weeks and invite the candidate to increase the dose or bring the injection forward. Both answers are wrong for the same reason, and the same confusion at the bedside escalates the dose inside the accumulation window, where the rise that follows would have happened anyway.

11.3 Reading a depot pharmacokinetic table without fusing its columns

Three quantities are tabulated for every preparation and they answer three different questions. Time to peak says when the concentration after a single injection is highest, and it is emphatically not the time at which a therapeutic concentration is first reached. Plasma half-life governs how fast the drug disappears once release has finished, and therefore how far an interval may be stretched. Time to steady state says when the patient's exposure stops climbing, and for a depot it is governed by the rate of release rather than by the half-life printed beside it. Candidates lose marks by fusing the first and the third.

PREPARATION	TIME TO PEAK (DAYS)	PLASMA HALF-LIFE (DAYS)	TIME TO STEADY STATE (WEEKS)
Haloperidol decanoate	3 to 9	21	about 14
Fluphenazine decanoate	8 to 12	7 to 10	about 8
Risperidone microspheres	about 28	3 to 6	about 8
Olanzapine pamoate	2 to 6	30	about 12
Aripiprazole monthly	4 deltoid, 5 to 7 gluteal	27	about 20
Paliperidone palmitate one-monthly	13	25 to 49	about 20
Paliperidone palmitate three-monthly	25 (disputed)	84 to 95 deltoid, 118 to 139 gluteal	about 52
Paliperidone palmitate six-monthly	29 and 32 by strength	148 and 159 by strength	not known

Read the columns as given and never derive one from another. Three half-lives of the tabulated plasma half-life of risperidone microspheres put the plateau inside a fortnight; the same row tabulates steady state at about 8 weeks, and that mismatch runs through most of the table. The half-life column describes elimination once release has finished; for a depot it is absorption from the injection site, not elimination, that decides when the concentration stops climbing.

Two things in that table repay attention. The first is fluphenazine decanoate, whose tabulated time to peak is not undisputed: the same footnote records estimates placing peak concentrations only a few hours after injection, and suggests the preparation likely produces two peaks rather than one. The second is the three-monthly paliperidone row, where the tabulated time to peak is contradicted by the prose of the same chapter of the same guideline, which gives a longer median with a stated range. The disagreement is internal to one book, so no candidate settles it by finding a better source. Declare both, and where a stem forces a single figure prefer the chapter's prose median, which is stated with a range, over the bare tabulated point, which is not; then move to the half-life, where the two do not differ.

MNEMONIC

Vehicle, Peak, Plateau, Cover

Four questions answer most long-acting injectable stems. **V**ehicle: oil, polymer microsphere or aqueous nanosuspension, which decides whether a test dose is wanted. **P**eak: when concentrations are highest after an injection, which is not when they become therapeutic. **P**lateau: when steady state arrives, which no dose and no interval can hurry. **C**over: whether the previous oral antipsychotic must continue, and for how long.

11.4 Risperidone microspheres: the release delay that has to be covered

Risperidone long-acting injection is the clearest demonstration in the class that release rate decides everything. It is not an esterified form of the parent drug but risperidone coated in polymer, and the microspheres must be suspended in an aqueous base immediately before use and stored in a fridge until then, a real constraint on community administration. Because the nature of the suspension requires the whole vial to be used, the available strengths of 25 mg, 37.5 mg and 50 mg for maintenance antipsychotic treatment are the only doses obtainable, and dosing flexibility is correspondingly poor. A test dose is neither required nor sensible.

The kinetics are the striking part. Peak concentration arrives at about 28 days while the plasma half-life is only 3 to 6 days. A preparation whose time to peak so far exceeds its half-life is describing release from the polymer, not elimination of the drug, and the practical translation is unforgiving: it takes 3 to 4 weeks for the first injection to produce therapeutic plasma levels, so the previous antipsychotic must be maintained at full dose for at least 3 weeks after the first injection, and oral cover is sometimes needed for as long as 6 to 8 weeks.

The interval is fixed rather than chosen. Risperidone depot must be administered **every 2 weeks**, and the Product Licence does not permit a longer gap between doses, which makes it the least elastic of the preparations tabulated above.

GOOD TO REMEMBER

For the first month of risperidone long-acting injection the oral antipsychotic is the treatment and the injection is not. Stopping oral cover on the day of the first injection, which feels tidy, removes the only active drug the patient has. Write the cover period into the prescription with an end date, and review that date rather than assuming somebody else will.

PITFALL

Reaching for risperidone microspheres in an acutely unwell patient who is refusing oral medication. It is intuitive and wrong, because the first injection releases too little drug to treat anyone for the best part of a month and there is, by definition, nothing available to cover the gap. A preparation with a delayed release is a maintenance decision, never an acute one.

11.5 Paliperidone palmitate: an ester prodrug and a deltoid loading regimen

Paliperidone is 9-hydroxyrisperidone, the major active metabolite of risperidone, and paliperidone palmitate is its ester prodrug. The ester is an aqueous nanosuspension that is hydrolysed to paliperidone after intramuscular administration and then slowly absorbed into the circulation. That chemistry is why it behaves unlike the risperidone microspheres from which it is chemically descended.

The initiation regimen makes the pharmacology visible. Paliperidone palmitate one-monthly is started, for maintenance antipsychotic treatment owned by the Schizophrenia: management, antipsychotics and adverse effects notes, with 150 mg intramuscularly into the deltoid, then 100 mg intramuscularly into the deltoid a week later within a window of plus or minus 4 days; the same notes set maintenance thereafter at 50 to 150 mg each month into either the deltoid or the

gluteal site, within a window of plus or minus 7 days. The site is not arbitrary. A single intramuscular dose into the deltoid produces on average a **28 per cent** higher peak concentration than the same dose into the gluteal muscle, which is precisely why both loading injections are deltoid: they exist to reach a therapeutic concentration quickly. Improvement in psychotic symptoms has been observed as early as the fourth day, and oral supplementation during initiation is not required from a pharmacokinetic point of view.

Set that against the preceding preparation and the contrast is the whole teaching point of switching: two second-generation injectables, one of them the active metabolite of the other, and one needs three weeks of full-dose oral cover while the other needs none. Nothing in the receptor pharmacology predicts that difference; only the formulation does.

11.6 Stretching the interval: the three-monthly and six-monthly arithmetic

The longer paliperidone formulations are not new drugs and they are not new doses. They are the same total exposure delivered less often, and each one is gated on stability at the shorter interval. The three-monthly formulation is indicated for patients who are clinically stable on the one-monthly formulation and require no dose adjustment, and it is recommended that they have been treated for 4 months or more with the monthly preparation and that the last two monthly doses were the same. The prerequisite is not bureaucratic: an interval this long removes any opportunity to correct a dose quickly, so the dose carried across must already be known to be right.

ONE-MONTHLY DOSE	EQUIVALENT THREE-MONTHLY DOSE
50 mg	175 mg
75 mg	263 mg
100 mg	350 mg
150 mg	525 mg

The conversion is the examinable part, because it is not the one the interval suggests. Tripling the interval does not mean tripling the dose: each three-monthly antipsychotic dose is **3.5 times** the monthly dose it replaces. The site effect persists but shrinks, injection into the deltoid producing an average increase of 11 to 12 per cent in maximum plasma concentration compared with the gluteal route, well below the advantage seen with the monthly preparation, and this formulation is licensed for both sites rather than restricted at initiation. Its half-life differs by site as the table above shows, and it reaches steady state a full year after it is started, later than anything else tabulated there.

The six-monthly formulation narrows eligibility further. It is indicated for patients already maintained on 100 mg or 150 mg of the one-monthly formulation, ideally for 4 months or more with the same dose for the last two injections, or for those who have had at least one injection of 350 mg or 525 mg of the three-monthly formulation and need no dose change; it is given once every 6 months into the gluteal muscle as maintenance antipsychotic treatment and must not be administered by any other route. Its kinetics are stated by strength, the prescribing not taught

here: peak concentration at 29 days for the 700 mg antipsychotic dose and 32 days for the 1000 mg dose, with plasma half-lives of 148 days and 159 days respectively. The switch itself may be made within a window of plus or minus a week of a monthly due date and within a fortnight of a three-monthly one.

Its time to steady state is recorded as not known, and that absence is not a curiosity. Without the one figure a prescriber would actually use, the honest teaching is the consequence rather than the number: with a preparation this long-acting, any dose change may take several months to declare its clinical effect, because the drug must accumulate to a new steady state before the change can be judged.

One caution about the conversion arithmetic here. The guideline is not internally consistent about which doses have a six-monthly equivalent. Its switching table and its own eligibility sentence agree with each other, and a separate sentence in the same section excludes a three-monthly strength that both of them plainly include; that stray exclusion should be read as a typesetting corruption of the lower strengths. In the same table the row label for the higher six-monthly strength names a monthly dose the section elsewhere says has no six-monthly equivalent at all, and the dose that genuinely converts to it is 150 mg monthly, the higher of the two eligible maintenance doses of this antipsychotic. This is the single place in the paliperidone arithmetic where the source is not clean, and a candidate who has read only one of the two statements will answer confidently and wrongly.

11.7 Why the licensed interval is not negotiable in either direction

The commonest request in a depot clinic is to bring the next injection forward, and the pharmacology declines it. Every long-acting injectable antipsychotic can be given safely at its licensed dosing interval, bearing in mind the maximum recommended single dose, and there is no evidence that shortening the interval improves efficacy. The reason lies in the shape of the curve once accumulation is complete: at steady state, **trough plasma levels** immediately before and after a dose usually sit substantially above the threshold concentration required for a therapeutic effect. In most patients there is no trough to rescue. Less frequent administration is in any case preferable, since the injection site itself is a source of discomfort.

The patient who reliably deteriorates in the days before an injection falls due is therefore reporting something the kinetics do not support, and the timing of the recovery gives it away: plasma concentrations continue to fall, slowly, for some hours and with certain preparations for some days after an injection is given. An improvement that appears within hours of the needle cannot be a pharmacological one.

Drift in the other direction is a genuine pharmacological problem. Each missed or delayed injection subtracts real exposure across a year, and the evidence tracks the annual count rather than any individual late dose.

TRAP **Reciting a single annual injection count as the point at which monthly paliperidone stops working.** The same edition of the same guideline puts the cut in two places, one dose apart. Writing on this drug, it states that patients receiving fewer than 12 injections a year have an increased risk of relapse, which counts the eleventh injection as already inadequate. In its general chapter on long-acting injectables, two United Kingdom studies showed that patients receiving 10 doses a year or fewer were at substantially higher risk of relapse than those receiving 11 or 12, which counts the eleventh as safe. Answer the direction, on which both agree: relapse risk rises as the annual count falls short of a full set of monthly injections. Where a stem forces one number, prefer the study-anchored statement because it names its evidence, and say that the same book sets the bar higher elsewhere.

11.8 A hazard that belongs to the formulation, not the molecule

Olanzapine pamoate carries a risk that has nothing to do with olanzapine's receptor profile and everything to do with its salt. **Post-injection delirium sedation syndrome** is thought to occur when the pamoate salt is inadvertently exposed to a large volume of blood or plasma, so that it dissolves far more rapidly than intended and releases a large amount of olanzapine into the circulation at once; plasma levels may then exceed 800 micrograms/L, and confusion, delirium and somnolence follow.

The epidemiology shapes practice. The incidence is less than 0.1 per cent of injections, and 86 per cent of reactions occur within 1 hour of injection, with a mean time of 30 minutes, reactions fully resolving within 72 hours. One study placed the incidence lower still, at 0.044 per cent of injections, fewer than 1 in 2,000, with 91 per cent of reactions apparent within 1 hour. That distribution, and not the severity of any individual episode, is what justifies a post-injection observation period, mandated at at least 3 hours in the European Union and the United Kingdom.

An aqueous second-generation depot is not simply a safer version of an oil-based one; each vehicle brings a hazard of its own, the oil an allergic one that the test dose exists to detect, the pamoate suspension one that only observation detects. The acute management of this and other psychotropic emergencies belongs to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes, and the choice of a long-acting injectable for a particular patient belongs to the Schizophrenia: management, antipsychotics and adverse effects notes. What stays here is the kinetic reasoning: know the vehicle, know when the concentration peaks, know when it plateaus, and cover the gap between the two.

12 · Choosing between antipsychotics: the comparative profile

Within the antipsychotics, the differences a patient actually feels are rarely differences in efficacy. They are differences in what else the molecule binds. Two agents can hold the same antipsychotic action and produce entirely unlike consultations, because one also sits on histamine receptors and the other does not, or because one raises prolactin and the other lowers it. The comparative profile states those differences precisely: what separates one agent from another once the shared antipsychotic action is taken as given, and how a prescriber derives the

separation instead of memorising it. Examination questions cluster here because the ground rewards reasoning from a binding profile and punishes a learned grid.

Three separate comparisons live under this heading, and conflating them is the commonest way the topic goes wrong: adverse-effect liability, predicted from receptor pharmacology; comparative efficacy, which comes from trial evidence and carries its biases; and dose equivalence, which is arithmetic and the least trustworthy of the three. Which agent to start in a person with schizophrenia, how to sequence trials and what to do when they fail belong to the Schizophrenia: management, antipsychotics and adverse effects notes. The monitoring schedules that follow from a cardiometabolic or cardiac-conduction liability belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes, and antipsychotic selection in Parkinson disease psychosis to the Neuropsychiatry of systemic and neurological disease notes. What belongs here is the pharmacology those decisions rest on.

5	3	32
RECEPTOR TARGETS THAT GENERATE THE CLASS SIDE-EFFECT SIGNATURE	NEWER AGENTS THAT OUT-RAISE THE OLDER DRUGS ON PROLACTIN	ANTIPSYCHOTICS IN THE EFFICACY NETWORK META-ANALYSIS

12.1 Liability is predicted, not memorised

The single most useful habit in this topic is refusing to learn the comparative side-effect table as a table. A large grid of agents against adverse effects is unlearnable, decays within weeks and gives nothing back when a drug outside it appears. What generates the grid is a short mapping from receptor to complaint, and that mapping stays.

That mapping is not this section's to build; built twice, it becomes two maps. The prediction running from a receptor action to the complaint the patient reports has its own section in these notes, which walks it receptor family by receptor family, gives the joint account of weight gain, and warns that sedation reaches the map by more than one route. The receptor systems underneath it belong to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

What the antipsychotic class adds is that the mapping is dense enough here to do most of the comparative work on its own. Knowing which of those targets a molecule touches tells you most of what the patient will complain of early in treatment and, more usefully for this section, which axis a comparison between two agents should be argued on. Everything past it is the step the map cannot take: putting agents in order.

RULE **Read liability off the receptor profile; read magnitude off a published ordering.** The receptor mapping tells you which complaint an agent will generate, and therefore which axis to worry about. It does not tell you how much, because affinity differs between molecules sharing a target. Magnitude must come from a source that states an ordering on that axis; where none does, the honest answer names the axis and stops. Every error here is either a mapping used for magnitude or a magnitude quoted with no ordering behind it.

MNEMONIC

Three rulers: liability, efficacy, equivalence

Liability is predicted from receptor pharmacology and then ranked one axis at a time. **E**fficacy comes from trial evidence and carries the biases of that evidence. **E**quivalence is arithmetic assembled from pooled judgement and is the least trustworthy of the three. Answers go wrong when a claim built on one ruler is defended with another: a receptor prediction offered as evidence of efficacy, or an equivalent dose offered as a tolerability guarantee. Name the ruler before answering and the three questions stop contaminating one another.

12.2 The axes, and which sources actually state an ordering

Once the mapping is in place, comparison proceeds one axis at a time. The temptation is always to ask which agent is best overall. There is no such question. An agent at the favourable end of one ordering routinely sits at the unfavourable end of another, and a comparison that has not named its axis has not been made.

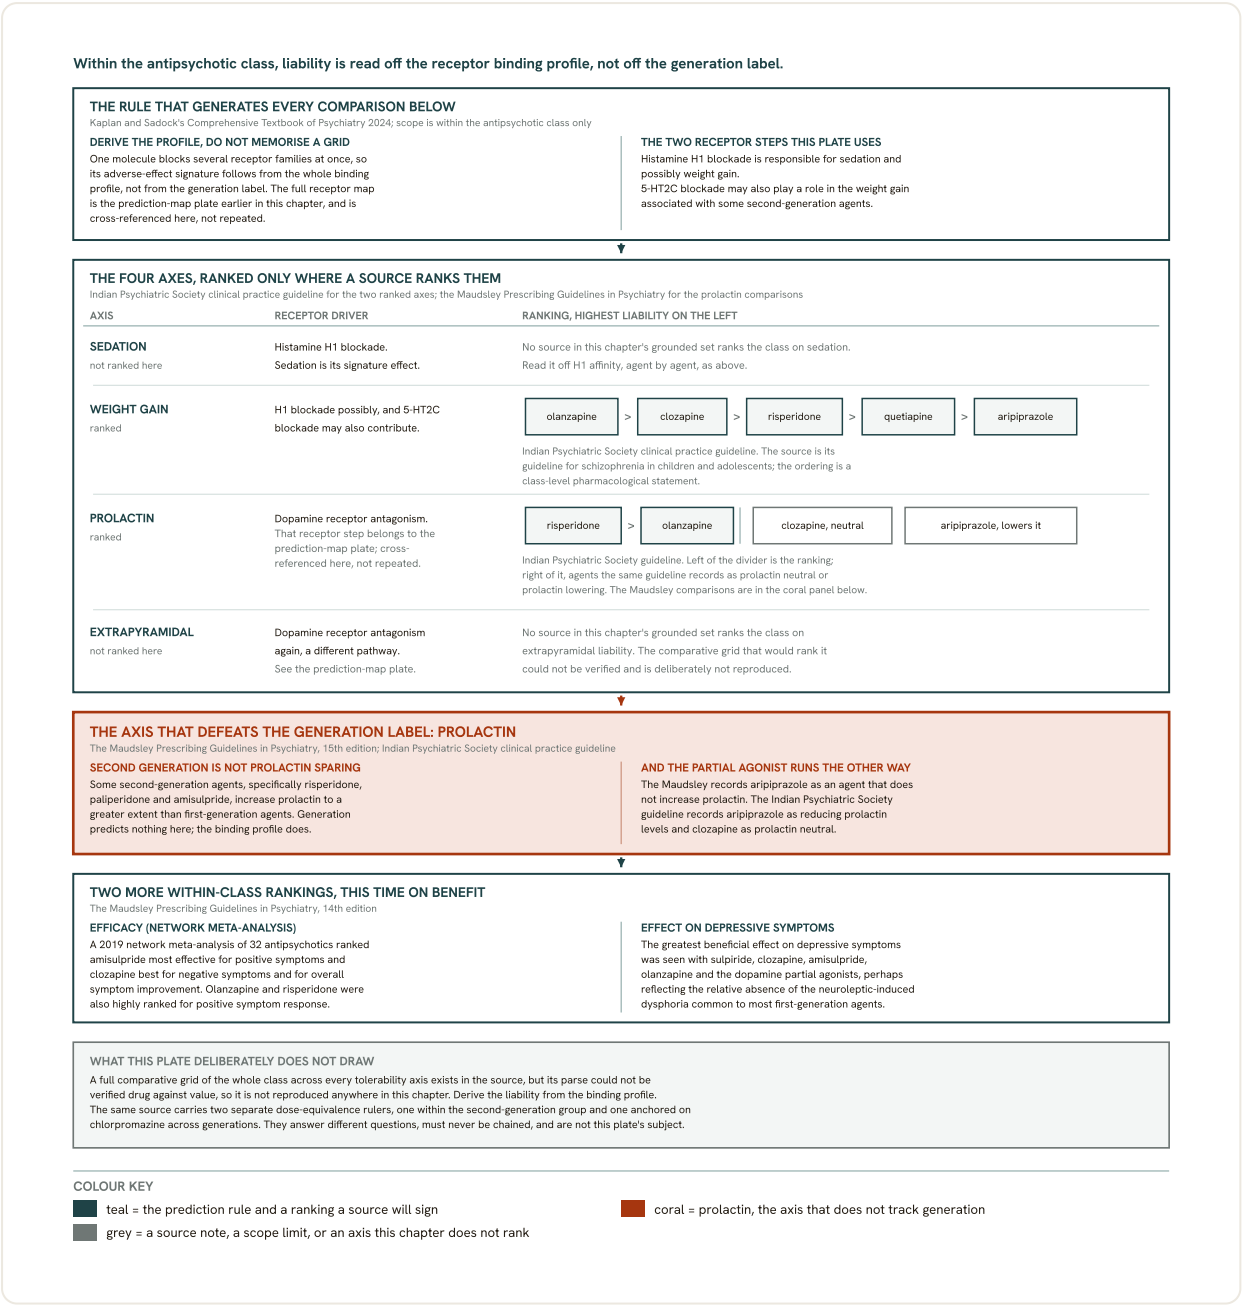


Fig 36.6 Within the antipsychotic class: what the receptor profile predicts, and where a source will actually rank the agents. Weight gain and prolactin are ranked where the Indian Psychiatric Society guideline ranks them, while sedation and extrapyramidal liability are drawn as unranked, because no source in this chapter's grounded set ranks them and the comparative grid that would has an unverifiable parse. The prolactin row is the one that defeats the generation heuristic: risperidone, paliperidone and amisulpride raise prolactin more than first-generation agents, while aripiprazole does not. (Indian Psychiatric Society Clinical Practice Guidelines for the Management of Schizophrenia in Children and Adolescents; The Maudsley Prescribing Guidelines in Psychiatry; Kaplan and Sadock's Comprehensive Textbook of Psychiatry 2024.)

The accompanying figure collects the within-class comparison. The table below states what is on record, axis by axis, with the body that states it. Read its gaps as carefully as its entries: an axis without a published ordering is not an axis on which any agent may be declared superior.

AXIS	THE ORDERING THE SOURCE STATES	SOURCE
Weight gain	Olanzapine greatest, then clozapine, then risperidone, then quetiapine, with aripiprazole least	Indian Psychiatric Society
Hyperprolactinaemia	Risperidone above olanzapine; clozapine prolactin neutral; aripiprazole reduces prolactin levels	Indian Psychiatric Society
Prolactin against the older agents	Risperidone, paliperidone and amisulpride raise prolactin more than the first-generation drugs	Maudsley Prescribing Guidelines
Positive symptoms	Amisulpride the most effective, with olanzapine and risperidone also highly ranked	Maudsley
Negative symptoms and overall improvement	Clozapine the best on both	Maudsley
Depressive symptoms	Greatest benefit with sulpiride, clozapine, amisulpride, olanzapine and the dopamine partial agonists	Maudsley

IN INDIA

The within-class orderings for weight gain and for hyperprolactinaemia quoted above are the ones issued by the Indian Psychiatric Society, and they are worth preferring over remembered international rankings for exactly that reason. Two cautions travel with them. They are published inside a guideline written for schizophrenia in children and adolescents, so treat the ordering as a class-level pharmacological statement rather than a licence for any population, and go to the paediatric guidance itself before acting on it in a young person. The guideline's own table also states the hyperprolactinaemia axis twice, with slightly different member lists, so the ordering is safest used at its ends, where the two statements agree, rather than for fine discrimination in the middle.

12.3 Sedation and weight: the histaminergic axis

Sedation and weight gain are grouped here because both carry a histaminergic contribution and are then dissociated by everything else. An agent with substantial affinity there tends to sit at the unfavourable end of both orderings, and the sedating agents of the class are broadly the weight-gaining ones. Neither complaint is a single-receptor effect, though. Central muscarinic blockade is a second route to sedation, so a strongly antimuscarinic agent is sedating for a reason the histamine column does not show and its position on a sedation axis cannot be read off histamine affinity alone; weight gain is taught in the section that owns the map as a joint action rather than as the property of either receptor. The correlation between the two orderings is real and loose, which is why the weight-gain ordering must be read as its own statement rather than inferred from how drowsy a drug makes people.

The ordering itself repays attention for a reason candidates often miss: it is not aligned with generation, with structural family or with any of the familiar labels. It runs across the modern agents, placing one of them at the top and another at the bottom. The axis discriminates within the newer drugs far more sharply than it discriminates between the eras, so an answer that reaches for the generational label to predict weight has reached for a variable with almost no explanatory power. What the labels do and do not encode is taught elsewhere in these notes.

12.4 Prolactin: the axis that ignores generation entirely

Prolactin is the axis where the intuitive ordering is not merely imprecise but inverted, which is why examiners like it best. The Maudsley Prescribing Guidelines record that risperidone, paliperidone and amisulpride raise prolactin by more than the first-generation agents do, placing three thoroughly modern drugs below the drugs they were expected to improve upon. The Indian Psychiatric Society ordering runs the same way within the newer group and adds two agents at the other end: clozapine is recorded as **prolactin neutral**, and aripiprazole as reducing prolactin levels rather than merely sparing it.

Aripiprazole is worth a sentence of its own because two independent sources converge on it from different directions. The Indian guideline places it below neutral on the endocrine axis, and the Maudsley, in the comments column of its own dosing table for the long-acting injectable preparation, records aripiprazole as **an agent that does not increase prolactin**. Convergence of that kind is the strongest evidence a within-class ordering ever offers, because the two statements rest on different bodies of data. The receptor theory that explains why a partial agonist behaves this way is taught elsewhere in these notes; what the comparative profile contributes is only that the behaviour is real and is recorded twice.

Note what these orderings are not. They rank the tendency to raise a hormone, and say nothing about the threshold at which a raised prolactin becomes clinically significant, how often it should be measured, or the monitoring that follows across the psychotropics as a whole, all of which belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

12.5 The autonomic axes: alpha and muscarinic

The remaining receptor-predicted effects attract less examination attention and more clinical trouble. Alpha blockade splits into two quite different consequences: **orthostatic hypotension** from the alpha-1 target, and a contribution to **sexual dysfunction** from the alpha-2 target, which Kaplan and Sadock records for this class and which the general map does not carry. Both are commonly misattributed, the first to sedation and the second to the illness or to the patient's mood. Muscarinic blockade produces a signature rather than a symptom, and it is the tightest of all the receptor-to-complaint predictions: **dry mouth, blurred vision, constipation and urinary retention** arrive together because they are one pharmacological event expressed in four organs.

The clinical value of a signature is that it can be recognised whole. A patient reporting any one of those four should be asked about the other three; a positive answer identifies the mechanism

immediately, and with it the axis on which the agent should be compared with its alternatives. This is the receptor mapping used in the direction it is best at: from complaint back to receptor, and from receptor to the comparative question worth asking.

PITFALL

Treating the anticholinergic cluster as four independent problems. A patient prescribed an artificial saliva, a laxative and reading glasses, then referred for urinary hesitancy, has been managed four times for one receptor event. The same error runs the other way with the alpha effects: postural dizziness recorded as drowsiness sends the reasoning to histamine when the target is alpha-1, and the comparative question is then asked on the wrong axis. Take the anticholinergic symptoms as a set, separate dizziness on standing from sleepiness at rest, and the receptor is usually obvious.

12.6 Comparative efficacy, and what a ranking encodes

Efficacy comparison rests on different evidence and behaves differently. A **network meta-analysis** of 32 antipsychotics published in 2019 and quoted by the Maudsley ranked amisulpride as the most effective drug for positive symptoms and clozapine as the best for both negative symptoms and overall symptom improvement, with olanzapine and risperidone also highly ranked for positive symptom response. Two features of that result deserve more attention than the ordering itself.

The first is that the ordering is a ranking, not a set of distances. Neighbouring positions in such a sequence may be separated by a difference no patient could detect, while distant positions are separated by a great deal. Quoting a rank without any sense of the spread is how a modest average difference becomes, in a candidate's answer, a categorical claim about superiority.

The second is a systematic distortion the source itself declares. More recently introduced drugs tend to show lower estimated efficacy because placebo response has been rising steadily since 1970. A drug tested against a more responsive placebo arm looks weaker without being weaker, so the ordering partly encodes when each agent happened to be trialled. Being able to say whether you are quoting a pharmacological fact or an artefact of the evidence base is most of what separates a good answer here from a merely correct one.

TRAP **Reading the efficacy ranking as evidence that the newer molecules are weaker.** The source publishing the ranking also states that estimated efficacy falls with year of introduction because placebo response has risen, which is a property of trials and not of drugs. Candidates quote the ordering, omit the caveat, and then defend a conclusion about molecules using data about study populations. The safe form of the answer states the ranking, names the drift, and declines to convert either into a claim that one agent is pharmacologically stronger than another.

12.7 Depressive symptoms as a comparative axis

The class also differs in what it does to mood, and the pattern is instructive because it is explained by an absence rather than by an action. The greatest beneficial effect on depressive symptoms is seen with sulpiride, clozapine, amisulpride, olanzapine and the dopamine partial

agonists, and the explanation the source offers is the relative absence of the neuroleptic-induced dysphoria common to most first-generation agents.

That is a genuinely different kind of comparative statement. The agents named do not appear to be antidepressant; they appear to be less dysphoriant. **Neuroleptic-induced dysphoria** is a drug effect, and an agent that produces less of it will show a larger improvement in depressive ratings without having acted on depression at all. This is subtracted harm masquerading as benefit, and it recurs throughout comparative psychopharmacology. It also predicts the shape of the list: the agents named sit at the gentler end of dopamine antagonism, and the partial agonists are on it for the same reason the potent antagonists are not. The consequences of dopamine occupancy across the pathways are taught elsewhere in these notes.

12.8 Dose equivalence: two rulers, and why they never compose

The last comparison is the least reliable and the most often misused. Equivalence tables exist because a prescriber switching agents needs some notion of comparable exposure, and they are assembled from pooled clinical judgement and trial dosing rather than measured. The Maudsley Prescribing Guidelines publish two, and they answer different questions.

Within the newer group, the guideline treats amisulpride 400 mg (the Schizophrenia: management, antipsychotics and adverse effects notes own the prescribing), aripiprazole 15 mg, **olanzapine 10 mg** (not taught here), quetiapine 400 mg and risperidone 4 mg as **approximately equivalent daily doses** of a second-generation antipsychotic, an equivalence acted on elsewhere. Across the generations the arithmetic differs: the older table is anchored on **chlorpromazine 100 mg/day**, an anchor owned by the clinical chapters and not by this one, and the same chapter states that, very approximately, that dose of antipsychotic is equivalent to **risperidone 1.5 mg**. Both are the source's own rough guides, and the second is hedged further still.

- The equivalences are a rough guide only, and the source records considerable disagreement between references about the exact figures.
- Clozapine is excluded from the table altogether, so no equivalent exposure can be read off it for that agent.
- The cross-generation figure carries its own additional warning, because potency comparisons that cross the generational boundary are less secure than comparisons inside one group.
- Equivalence is a statement about dose and about nothing else. Two doses that the table calls equivalent may sit at opposite ends of the weight-gain ordering, so an equivalent switch is not a tolerability-neutral switch.

TRAP **Chaining the two equivalence tables to derive a chlorpromazine equivalent for a newer agent.** The within-class table and the cross-generation figure are separate rulers, built by different methods to answer different questions, and each is called a rough guide by the source that publishes it. Composing them yields a number that looks authoritative, has no source behind it, and will be defended in a viva precisely because the arithmetic felt sound. This is a different error from misreading an efficacy ranking: that one over-reads real data, this one manufactures data never published. State both frames, and use each only inside the comparison it was built for.

GOOD TO REMEMBER

When a switch is being planned, the equivalence table answers only where to land, never how the patient will feel on landing. Look up the target dose, then look up the new agent separately on each liability axis that matters for this person, because the two lookups are independent and the table encodes nothing about the second. When the drug in question is one the table does not cover, whether because its equivalence was unknown or because it is excluded by design, say so rather than interpolate.

Assembled, the topic reduces to a sequence. Identify the axis the comparison is about; predict from the receptor mapping which agents are at risk on it; then look for a published ordering, preferring one issued where you practise. If an ordering exists, quote it with its scope; if none exists, name the axis and decline to rank.

Liability is predicted from what else the molecule blocks and ranked one axis at a time against a source that actually states an ordering; the generational label ranks nothing; and the two dose-equivalence tables are separate rulers that must never be chained.

Antidepressants

13 · The antidepressant families: one mechanism axis, five groups

Learn the antidepressants by where the molecule acts, not by the decade it was launched in. Every drug in this class raises a monoamine's availability at the synapse, and there are few places in the synaptic machinery where that can be done: at the transporter clearing the transmitter out of the cleft, at the receptors the transmitter acts on, or at the enzyme that destroys it. Each position is a landing point, each generates its own burden of unwanted effects, and the families an examiner asks about are the drugs sharing one. Put the axis in place first and the receptor lists become derivable rather than memorised.

The marks here are not for family names but for the discriminations the axis makes visible: why an agent blocking one transporter is tolerated quite differently from one blocking several receptor families besides; why the group that displaced the older drugs displaced them on safety rather than efficacy; and why every family, whatever its target, keeps the patient waiting the same length of time. Choosing an antidepressant for a particular patient with a depressive episode

belongs to the Mood disorders: pharmacotherapy notes, and the transmitter systems as biochemistry belong to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

1957	1958	several weeks
THE MONOAMINE OXIDASE INHIBITORS ENTER PRACTICE	THE TRICYCLICS FOLLOW, FROM AN OPPOSITE MECHANISM	TO A THERAPEUTIC EFFECT, WHATEVER THE FAMILY

13.1 One axis, and a family count the sources do not fix

It is tempting to open with a numbered list of families, and the references will not support one: they disagree on how to cut the class. KDT Essentials of Medical Pharmacology says outright that a cogent classification of the antidepressants is difficult and offers a working classification rather than a definitive one; other texts group by generation, and others by first-line status rather than mechanism. The count is an artefact of which book is open. Say that the sources differ, then teach what does not differ.

What survives across all of them is the axis, because the axis is a statement about molecular targets rather than nomenclature. A monoamine's concentration in the cleft is set by how fast it is released, how fast it is taken back and how fast it is broken down, and a drug can intervene at only one of those points. A molecule that also binds receptors directly adds actions with nothing to do with raising transmitter availability at all.

That distinction separates therapeutic action from adverse-effect signature more cleanly than any generational label does: the transporter and the enzyme are where the treatment happens, and the receptors a molecule touches on the way past are where the complaints come from. Every family below answers the same pair of questions: what does this drug block, and what else does it block.

RULE **An antidepressant family is a position on one mechanism axis, defined by the molecular target its drugs share; the number of families is a property of the textbook, not of the pharmacology.** Teach the axis and the family names arrive as consequences. Teach the names first and every source that cuts the class differently becomes a contradiction.

MNEMONIC

Transporter, Receptor, Enzyme, Exception

The landing points, in the order they are usually taught. **T**ransporter, for the selective reuptake inhibitors, whose whole definition is one shared transporter. **R**eceptor, for the older agents adding direct receptor blockade on top of it. **E**nzyme, for the classical monoamine oxidase inhibitors, which touch no transporter. **E**xception, for the agent whose reuptake block reaches the dopamine transporter. Recited so, it delivers the axis rather than a count, which is all the sources support.

13.2 Two founding families, mechanistically opposite, a year apart

The chronology explains why the class looks as it does. The Companion to Psychiatric Studies dates the monoamine oxidase inhibitors to **1957** and the tricyclics to **1958**. So the two mechanistically opposite founding families arrived within a year of one another, one disabling an

enzyme and the other blocking transporters, and neither was designed from a theory of what depression is.

That is not a footnote, it is the reason the monoamine account of mood disorder exists at all. The theory was reverse-engineered from drugs already observed to work, and built to explain what they had in common: each raises monoamine availability in the cleft, by opposite routes, so availability became the candidate explanation. Every family that followed was designed to reproduce that feature with fewer of the other actions attached, which is why the later history of the class reads as a narrowing of pharmacology rather than new therapeutic ideas.

Two consequences follow. The axis describes what these drugs do and not what depression is, and a stem offering a receptor abnormality as the cause of an illness because a drug acting there relieves it is testing exactly that confusion. And the founding families sit at opposite ends, so everything between them is a compromise.

13.3 Transporter blockade: one shared target, differences that sit elsewhere

The most populated position on the axis is selective blockade of a single transporter, and its definition is exact. The selective serotonin reuptake inhibitors share one common mechanism of action, blockade of the **serotonin transporter**, and they differ in their actions at additional receptors. Kaplan and Sadock's Comprehensive Textbook of Psychiatry states the difference at those receptors; it does not say receptors are the only place members differ, and the within-class section of these notes adds clearance and enzyme inhibition.

That definition does more work than it looks. It says first that the therapeutic action is a family property: whatever a member does that the others do not, it is not doing it at the transporter they all share. It says second that differences between members are pharmacological rather than commercial, since a difference at an additional receptor is a difference in what the drug does to the patient. It says third where to look when a member behaves unexpectedly: at what else the molecule carries, not at the target they all share.

Shared target, shared effect; unshared receptor actions, unshared adverse effects. The serotonin and noradrenaline reuptake inhibitors that sit alongside this family, and the general business of reading a receptor-binding profile into a predicted complaint, are taken up elsewhere in these notes.

GOOD TO REMEMBER

When a patient reports that a drug from one of these families did not suit them, establish what actually happened before writing off the family. The members share blockade of one transporter; what differs between them sits elsewhere, at additional receptors and in how the drug is handled. An intolerance generated away from the shared target says nothing about that target and need not transfer to the next member; an effect following from the shared action does. Which agent should follow in a depressive episode belongs to the Mood disorders: pharmacotherapy notes.

13.4 Which transporter buys which property

Blocking a transporter is not one action, because there is more than one monoamine transporter and a molecule may engage any combination of them. KDT Essentials of Medical Pharmacology draws the mapping between transporter and clinical property, and draws it carefully, introducing it as a set of tentative conclusions rather than as settled pharmacology. The mapping is this: inhibition of noradrenaline and serotonin uptake is what correlates with antidepressant action, whereas **inhibition of dopamine uptake** correlates with stimulant action.

One sentence, and it accounts for a great deal that would otherwise be memorised drug by drug: why the antidepressant families were built around the noradrenaline and serotonin transporters; why the powerful dopamine uptake inhibitors are classified as stimulants and drugs of abuse rather than antidepressants, though they raise a monoamine in the cleft exactly as the antidepressants do; and why raising all three transmitters is not automatically better than raising two, since the third transporter brings a different clinical property rather than more of the same one.

The hedge is not decoration and should survive into the answer: the claim is a correlation offered tentatively, not a rule that dopamine uptake blockade can never be antidepressant. The source exempts an agent in the same breath, and the exemption is the instructive part.

That agent is bupropion, defined as a **dual noradrenaline and dopamine reuptake inhibitor**. It engages the transporter the mapping associates with stimulant action and is an antidepressant nonetheless. Kaplan and Sadock's Comprehensive Textbook of Psychiatry records that it is not often used as a first-line agent for anxiety disorders, which is legible rather than arbitrary: an agent whose distinguishing feature is a dopaminergic action is not the obvious choice where activation is unwelcome. Its full profile is taken up elsewhere in these notes.

TRAP **Hardening the transporter mapping into an exclusion: "dopamine reuptake inhibition is never antidepressant."** The source states a correlation, states it tentatively, and exempts bupropion explicitly in the same passage. The correlation is useful, because it explains why the stimulants are not antidepressants. Turned into a law it becomes false, and a stem naming a dual noradrenaline and dopamine reuptake inhibitor and asking for its class will catch anyone who converted a tendency into a prohibition.

13.5 Reuptake plus direct receptor action: the older signature

The next position is not a different target but an additional one, and it defines the oldest surviving family. KDT Essentials of Medical Pharmacology makes the point definitionally: the **first generation** of antidepressants is defined not by uptake blockade alone but by the fact that these compounds also have direct effects at adrenergic, cholinergic and histaminergic receptors, and that added receptor action is what separates them from everything that came after. The label is mechanistic and not merely chronological, because a generational name invites the assumption that it records a launch date.

Kaplan and Sadock's Comprehensive Textbook of Psychiatry states the same architecture from the tricyclic end and adds the consequence. These agents inhibit the reuptake of more than one

monoamine, noradrenaline and serotonin among them, and in addition block cholinergic, histaminic and **alpha-1 adrenergic** receptors, and it is that second set of actions, rather than the reuptake blockade, that generates their side-effect profile. Read that twice: the reuptake action and the receptor actions do different jobs. The reuptake action treats the illness; the receptor actions produce the complaints.

Once that split is in place, the adverse-effect burden of the older family stops looking arbitrary. It is not the price of a stronger antidepressant action; it is the price of an unrelated set of receptor actions carried by the same molecule. A drug reproducing the reuptake action without the receptor actions would keep the therapeutic effect and shed the burden, which is the design brief the later families were built to and the comparison the **second generation** label was coined to record. The tricyclic profile in full, including what receptor promiscuity does to the heart, is taken up elsewhere in these notes.

13.6 Enzyme inhibition: the far end of the axis

The last founding position attacks the problem from the opposite direction, and is the only family that touches no transporter at all. **Enzyme inhibition** works on the transmitter's destruction rather than on its clearance from the cleft: the classical monoamine oxidase inhibitors irreversibly inactivate both isoforms of the enzyme, and serotonin, noradrenaline and dopamine consequently rise together.

Three properties of that position matter, and all follow from the mechanism. Irreversible inactivation means the effect outlasts the drug, because what has to recover is not a falling plasma concentration but a regenerating population of enzyme. Non-selectivity between the isoforms makes this the one position on the axis raising all three monoamines by design rather than as a side issue. And acting on an enzyme means the drug alters the handling of everything else that enzyme processes, which is why this family is defined in practice by what cannot safely be taken alongside it.

The reversible inhibitors are the exception and behave differently for that reason. The class read as an interaction problem, with its dietary constraint and washout arithmetic, is taken up elsewhere in these notes; the emergencies those interactions can produce belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

POSITION ON THE AXIS	WHAT THE DRUG ENGAGES	THE FAMILY THERE	WHAT IT BRINGS WITH IT
Selective transporter blockade	Blockade of the serotonin transporter, shared by every member	The selective serotonin reuptake inhibitors	Members differ in their actions at additional receptors
Transporter plus direct receptor action	Reuptake of more than one monoamine, with cholinergic, histaminic and adrenergic receptor blockade added	The first generation, including the tricyclics	The added receptor actions, not the reuptake block, generate the side-effect profile
Enzyme inhibition	Irreversible inactivation of both isoforms of monoamine oxidase	The classical monoamine oxidase inhibitors	Serotonin, noradrenaline and dopamine rise together
Transporter blockade including dopamine	Dual noradrenaline and dopamine reuptake inhibition	Bupropion, the exception	Not often used as a first-line agent for anxiety disorders

13.7 What moving along the axis actually bought

If the later families were built to reproduce the therapeutic action without the receptor actions, the question is what that redesign delivered. The answer is the most examinable statement here. Kaplan and Sadock's Comprehensive Textbook of Psychiatry is unusually direct: a very important consideration in the adoption of the selective reuptake inhibitors as first-line treatments was their **wide therapeutic index**, and deaths from tricyclic overdose fell once those agents supplanted the older class as the mainstay of treatment. The same text states plainly that the selective agents have never been shown to be more effective than any other class of medication for depression.

So the axis, read from the promiscuous end to the selective end, is a safety gradient and not an efficacy gradient, and that deserves its mechanism. A drug confined to one transporter has little left to do at a toxic dose except more of what it already does. A drug additionally blocking several receptor families has, at a toxic dose, actions already unwanted at a therapeutic one, and those actions scale. Narrowing the pharmacology narrowed the overdose by the same route.

That is not an abstraction: self-poisoning is a foreseeable event in this population. What a therapeutic index is, and which psychotropics need a plasma level, is taken up elsewhere in these notes; the management of an overdose belongs to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

TRAP **Treating the arrival of the selective agents as an advance in efficacy.** It was an advance in therapeutic index. The source recording the adoption records in the same passage that these agents have never been shown to be more effective than any other class for depression, and the fall in deaths it reports is an overdose statistic rather than a response rate. Candidates conflate the two because the newer family became first-line, and first-line reads as better.

13.8 The one property the axis does not predict

Everything so far has been a case for the axis: the target predicts the therapeutic action, the receptors touched on the way predict the complaint, and the position predicts what happens at a toxic dose. There is one clinically decisive property it does not predict, and that property is the same wherever the drug sits. **Delayed onset of action** is a property of the class rather than of any one drug or family: the pharmacological action begins with the first dose, but the therapeutic effect does not, and antidepressants may take several weeks to produce it.

That gap is the strongest argument that transporter or enzyme blockade is the beginning of the therapeutic mechanism rather than the whole of it: whatever eventually lifts mood is downstream of the blockade, takes time to develop, and is not named by the axis.

Four consequences follow, and none needs a number.

- Adverse effects arising directly from receptor blockade are available from the first dose, because that action is immediate; the benefit is not on the same timetable.
- Early in treatment the patient carries the whole cost of the molecule's receptor profile and none of its return, which is the period in which treatment is most often abandoned.
- A drug's position on the axis predicts what the first fortnight feels like and very little about when the illness moves.
- Because the delay belongs to the class and not the family, a different landing point buys a different adverse-effect signature, not a faster answer.

PITFALL

Reassuring a patient that the early adverse effects of a receptor-promiscuous agent are a sign the drug is starting to work. They are the opposite: immediate receptor pharmacology arriving on schedule while the therapeutic action is still weeks away, and arising from the very actions that are not treating the illness at all. Patients given that reassurance stop the tablet in the window where it has delivered every cost and no benefit, and the discontinuation is recorded as intolerance.

Every antidepressant family is a position on one axis: blockade of the serotonin transporter alone, reuptake blockade of more than one monoamine with cholinergic, histaminic and adrenergic receptor blockade added, irreversible inactivation of both isoforms of monoamine oxidase, or a reuptake block that includes the dopamine transporter; the therapeutic action comes from the transporter or the enzyme, the adverse-effect burden comes from the receptors, and what the selective families bought was a wide therapeutic index and not greater efficacy.

14 · SSRIs compared: what actually differs inside the class

The selective serotonin reuptake inhibitors are the class you are most likely to be asked to compare with itself. Almost no marks are available for whether an SSRI works; they are available for the harder question of what actually separates one molecule in this class from another,

because a candidate who believes the agents are interchangeable will say so within a sentence, and a candidate who has read the pharmacology will not. The differences are real, they are few, and they sit in a short list of places: what the drug does at receptors other than its target, how long it takes to leave the body, and how heavily it interferes with the enzymes that clear everything else. Nearly everything else about these drugs is shared, and knowing which properties are shared is exactly as examinable as knowing which are not.

What follows treats the class as a pharmacological object rather than as a treatment. Choosing an antidepressant for a patient with a depressive episode, judging response and switching agents belong to the Mood disorders: pharmacotherapy notes. Prescribing for children and adolescents, and the paediatric suicidality warning, belong to the Child psychiatry: assessment, treatment, services and protection notes. Interactions across drug classes, prescribing where an organ has failed, and serotonin toxicity as an emergency belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

About 80%	Four to six days	4 to 16 days	Two to four days
STRIATAL TRANSPORTER OCCUPANCY AT MINIMUM THERAPEUTIC DOSES	FLUOXETINE HALF-LIFE ON CHRONIC DOSING	ITS ACTIVE METABOLITE, WHATEVER THE DURATION OF DOSING	BEFORE DISCONTINUATION SYMPTOMS APPEAR AFTER STOPPING

14.1 One target, and a threshold rather than a dose

Every drug in this class does the same primary thing: it blocks the presynaptic serotonin transporter and leaves the rest of the receptor landscape largely alone. The consequence that reframes the whole class is that the quantity which has to be achieved for an antidepressant effect is not a milligram figure at all, but a proportion of transporter occupied. Positron emission tomography performed on citalopram, fluoxetine, sertraline, paroxetine and extended-release venlafaxine, each given at its own minimum therapeutic dose, shows **about 80% occupancy** of serotonin reuptake sites in the striatum. Different molecules, different milligram figures, one shared **occupancy threshold**.

The consequences reach in opposite directions. Comparing agents in milligrams is meaningless, because a milligram figure is a property of the molecule, of its potency and of how the body handles it, and not of the target that has to be occupied; the sentence "sertraline is a weaker drug than paroxetine because the tablets are larger" is a pharmacological error rather than an approximation. Occupancy is also regional and not uniform, and that regional variation carries a teaching point of its own: occupancy in the orbitofrontal cortex runs lower than in the striatum, which is the stated explanation for obsessive-compulsive disorder requiring higher doses than depressive illness does. That is an observation about where the drug binds and how much of it binds there, not a statement about how any disorder is managed.

14.2 Receptor cleanliness, and the axes that survive it

Receptor cleanliness is the property that made this class replace its predecessor, and it is worth stating as a negative fact because that is how it behaves clinically. Most SSRIs carry essentially no clinically meaningful antagonism at muscarinic, histaminergic, adrenergic or postsynaptic

serotonin receptors, and that absence is the entire reason the class is tolerated where the tricyclics are not. The point is worth labouring: the burden that makes a tricyclic hard to take is not serotonergic at all. It is the sum of what that molecule does at every receptor it was never aimed at, and an agent that reaches the transporter without visiting those receptors inherits the benefit and not the cost. The mapping from a binding profile to the adverse effect the patient reports is set out in the section of these notes on reading a receptor profile, and the receptor promiscuity of the tricyclics themselves is taught with that class.

Cleanliness at receptors is not the same as sameness, and what survives it is the material this section exists to teach.

- **Off-target binding.** A minority of the class does visit other sites, and the sites visited predict the complaint the patient brings back.
- **Elimination kinetics.** Serum half-lives across the class run from hours to days, and the position an agent occupies on that spread predicts what stopping it will feel like.
- **Enzyme inhibition.** The agents differ in how strongly they inhibit the cytochrome P450 isoenzymes, and it is this, not the serotonergic action, that makes one SSRI a harder drug to add to an existing regimen than another. The map of the isoenzymes themselves, and of what varies in them, is set out in the section of these notes on the psychotropic substrate map; which SSRI inhibits which isoenzyme is not carried by the sources used here. The interaction matrix spanning drug classes belongs to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

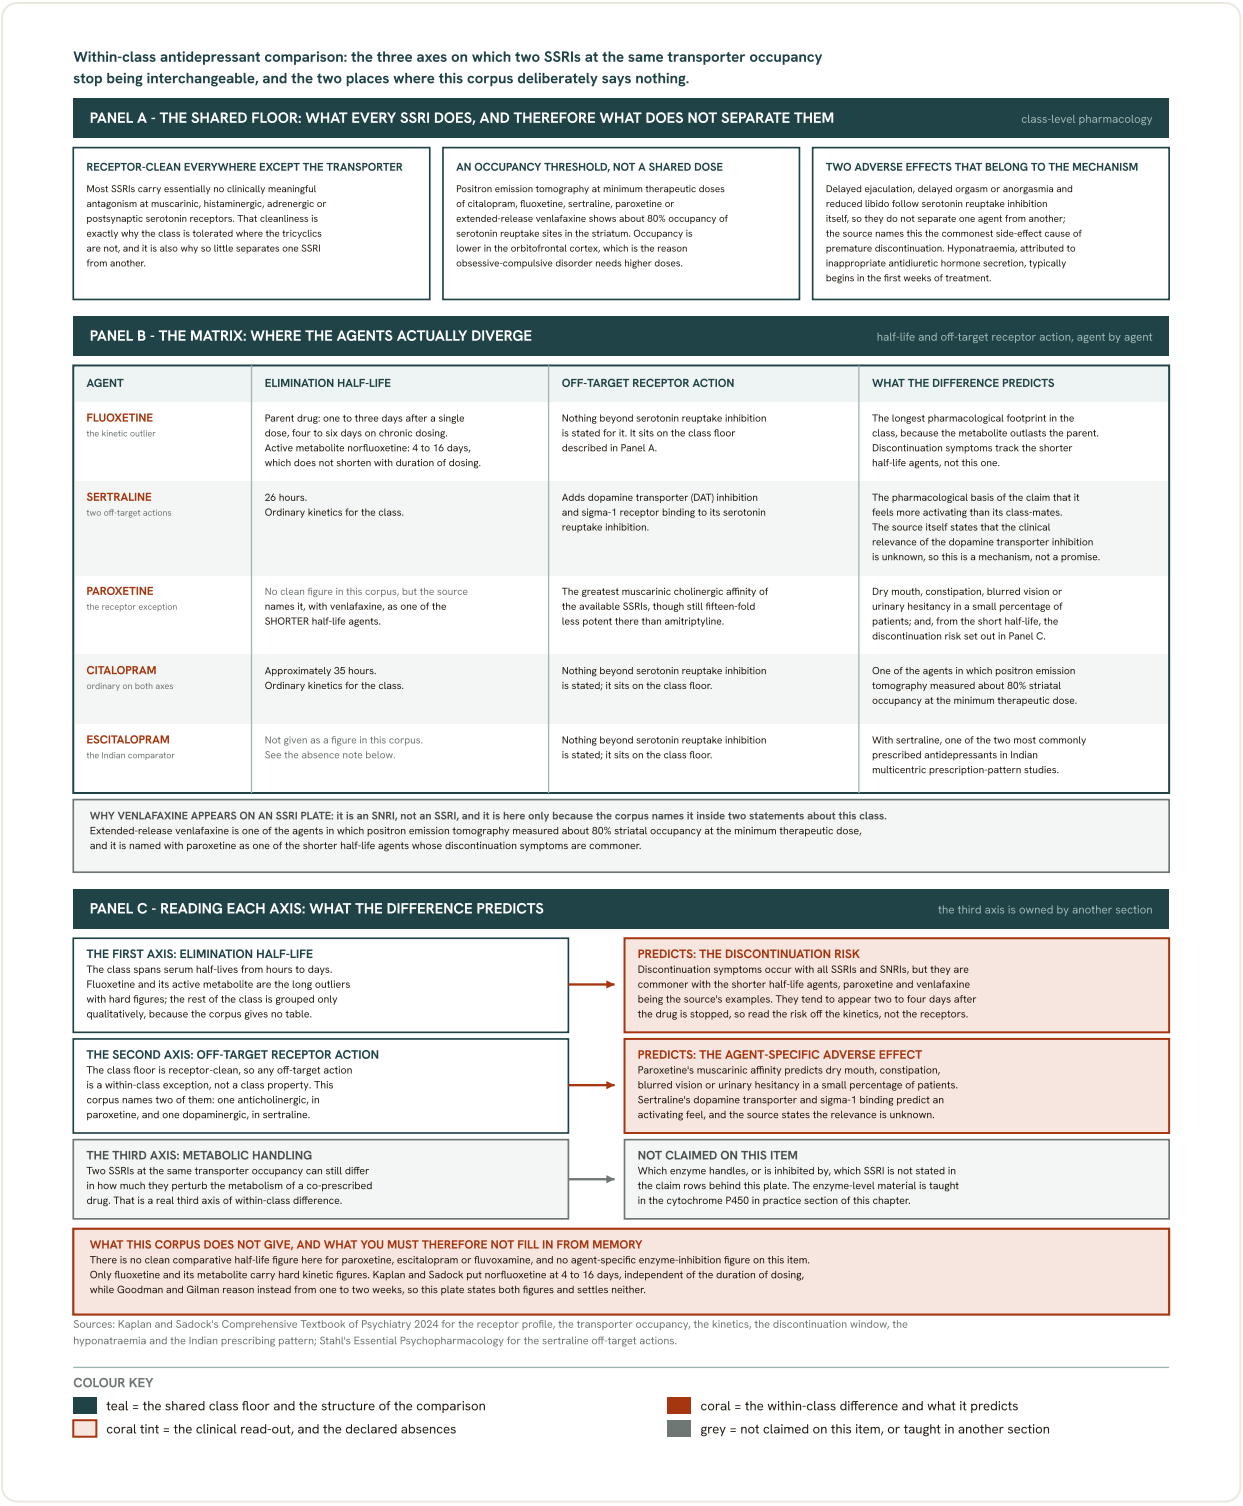


Fig 36.7 Within-class antidepressant comparison: half-life, off-target receptor action, and the axis this item does not claim. Most SSRIs are receptor-clean beyond the transporter, and every agent has to reach an occupancy threshold rather than a shared milligram dose, which is why sexual dysfunction and early hyponatraemia belong to the mechanism and not to any one drug. What separates the agents is kinetic and receptor-level: fluoxetine and its active metabolite carry the long half-lives, sertraline adds dopamine transporter and sigma-1 binding, paroxetine carries the greatest muscarinic affinity in the class, and the shorter half-life agents carry the discontinuation risk. The third axis, metabolic handling at the cytochrome P450 enzymes, is a genuine source of within-class difference that this item does not claim and that the cytochrome P450 in practice section of this chapter teaches. (Kaplan and Sadock's Comprehensive Textbook of Psychiatry 2024; Stahl's Essential Psychopharmacology.)

The accompanying figure draws the comparison out. Read the class this way and the examination question stops being a memory test and becomes a reasoning one: given an agent and a

complaint, which axis is the complaint sitting on.

RULE **What the SSRIs share is the target; what separates them is what else they touch, and how long they take to leave.** Every within-class comparison in this class reduces to that sentence. Shared mechanism generates the effects the whole class produces, and it therefore cannot be escaped by moving between agents. Off-target binding and elimination kinetics generate the differences, and those are the only differences a change of agent can buy.

14.3 The agents that break the class rule

The exceptions to that cleanliness attach to named agents, and the examiners know which agents they are. The first difference is **muscarinic cholinergic affinity**. Kaplan and Sadock's Comprehensive Textbook of Psychiatry, 11th edition records that paroxetine has the greatest affinity for muscarinic receptors of the available SSRIs, and therefore produces dry mouth, constipation, blurred vision or urinary hesitancy in a small percentage of patients. Keep that fact in proportion. The proportion of patients affected is small, so the anticholinergic profile of paroxetine is a within-class difference and not a class-changing one. And the affinity, though highest in the class, is modest in absolute terms: the same source records paroxetine as fifteen-fold less potent at muscarinic receptors than amitriptyline. Paroxetine is the anticholinergic SSRI; it is not an anticholinergic drug in the sense a tricyclic is.

The second exception runs the other way. Stahl's Essential Psychopharmacology describes sertraline as adding **dopamine transporter inhibition** and sigma-1 receptor binding to its serotonin reuptake inhibition, which is the pharmacological basis usually offered for the impression that sertraline feels more activating than its class-mates. Handle this one honestly in an answer, because the source does: the same passage states that the clinical relevance of that dopamine transporter inhibition is unknown. A mechanism that is present at the molecule and unproven at the bedside is exactly the kind of claim a viva rewards you for labelling correctly. Say that the binding is documented, say what it might explain, and say that the link to the clinical impression is not established.

MNEMONIC

Paroxetine, Sertraline, Fluoxetine

These names carry almost every within-class difference worth printing. **Paroxetine** is the one with muscarinic company, and the one whose short kinetics make stopping it uncomfortable. **Sertraline** is the one with company at another transporter. **Fluoxetine** is the kinetic outlier, the drug that is still present long after the patient has stopped taking it. Everything else in the class is, for examination purposes, the class.

14.4 Half-life: the axis with the widest spread

The **elimination half-life** is where this class separates most dramatically, and one agent accounts for most of the separation. Fluoxetine has an elimination half-life of **one to three days after a single dose and four to six days on chronic dosing**. The source gives both figures and not the reason they differ, and the reason is worth getting the right way round: a half-life lengthens when

clearance falls or the volume of distribution widens, which is the rule the section of these notes on pharmacokinetics and half-life sets out. Accumulation is what a long half-life then produces, not what lengthened it. That figure alone would make fluoxetine unusual. What makes it exceptional is the metabolite.

Norfluoxetine is pharmacologically active, and its elimination half-life is longer than the parent compound's and does not shorten however long the drug has been taken. The clinical translation is that fluoxetine's pharmacological footprint outlasts fluoxetine. Weeks after the last tablet, a patient may still be carrying a meaningful serotonergic load, which is why every question about washout, about starting a drug that interacts, and about the interval before a monoamine oxidase inhibitor, is answered differently for this agent than for the rest of the class.

The same arithmetic settles a smaller question that comes up constantly. A missed dose of a short half-life agent is a real fall in exposure, because a meaningful fraction of what is circulating is gone before the next tablet is due. A missed dose of fluoxetine is close to arithmetically invisible, because the interval between doses is a small fraction of the time the drug needs to halve. The clinical texture that follows from this, the impression that some of these drugs are forgiving and others are not, is not a property of the receptor at all.

A recital of half-life figures for the remaining agents is worth less than the principle they illustrate, and the principle is short. Serum half-lives across the class span from hours to days; fluoxetine and its active metabolite are the long outliers, and they are the ones with firm figures attached; and the gap between fluoxetine and its class-mates is several-fold rather than an order of magnitude. An examiner asking you to rank the class by half-life wants the outlier named and the direction right, not a table recited.

TRAP **Quoting a single settled half-life for norfluoxetine.** The standard references do not agree, and a candidate who has met only one of them assumes the other is a misprint and defends the wrong thing. Kaplan and Sadock's Comprehensive Textbook of Psychiatry, 11th edition puts the elimination half-life of the active metabolite at **4 to 16 days** and states that this is independent of how long the drug has been given. Goodman and Gilman's The Pharmacological Basis of Therapeutics reasons instead from one to two weeks when it sets the interval that must elapse before a monoamine oxidase inhibitor may be started. Neither source is wrong and the ranges overlap: the wider figure admits more between-individual variation than the working number the washout arithmetic is built on. The safe answer gives the range, names whose figure it is, and does not manufacture agreement. The washout interval itself is taught with the monoamine oxidase inhibitors.

14.5 Discontinuation, read off the kinetics

Discontinuation symptoms are the clearest illustration in the whole class of a clinical phenomenon that is predicted by pharmacokinetics and by nothing else. They can occur with all SSRIs and SNRIs, but are more common with the shorter half-life agents, paroxetine and venlafaxine being the named examples. Nothing about a receptor profile predicts which agents will be difficult to stop; the elimination curve predicts it completely. The agent at the other end of

the spread is fluoxetine, whose long tail means the concentration falls slowly enough to taper the patient without anyone intending it.

The mechanism underneath is worth stating because it makes the rule reconstructible rather than memorised. Continuous occupancy of the transporter over weeks of treatment is met by adaptation in the systems downstream of it. When a rapidly cleared agent is stopped, occupancy falls faster than those adaptations can unwind, and the mismatch is what the patient experiences. When a slowly cleared agent is stopped, occupancy falls under cover of the same adaptations reversing, and there is no mismatch to feel. The half-life is therefore not a proxy for the risk; it is the risk, expressed as a rate.

The timing carries the same lesson. Symptoms tend to appear **two to four days** after the medication is stopped, which is an interval measured in the drug's own elimination window. That is the discriminator worth carrying: a relapse of the underlying illness does not respect the pharmacokinetics of the tablet, and a syndrome that arrives inside the clearance window and resolves when the drug is reinstated is a kinetic event. The clinical handling of that event in a patient being treated for a depressive episode belongs to the Mood disorders: pharmacotherapy notes; what belongs here is the reason the risk differs between agents at all.

14.6 The effects that belong to the mechanism, not the molecule

Some adverse effects of this class are not properties of any agent within it. They are properties of inhibiting serotonin reuptake, and they follow the mechanism wherever it goes. **Sexual dysfunction** is the important one. Delayed ejaculation, delayed orgasm or anorgasmia and reduced libido are a common effect of serotonin reuptake inhibition itself, which is why the same complaint recurs across agents that share nothing else except the target. It is also the commonest side-effect reason for premature discontinuation, which is the fact that makes this more than a footnote: the effect that most often ends treatment is the one least amenable to changing the tablet.

Hyponatraemia behaves as a class hazard in the same way. SSRIs raise the risk of hyponatraemia, and it typically begins during the first weeks of treatment; the source offers inappropriate antidiuretic hormone secretion as the probable mechanism rather than a settled one, and an answer that hardens it into the mechanism has said more than the text does. Of the pair, the timing is both the firmer statement and the more useful one, because it tells you which stretch of treatment the explanation is available for. The monitoring schedule that follows from this across the psychotropics belongs to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

PITFALL

Reading sexual dysfunction as a property of the particular molecule. It is a property of serotonin reuptake inhibition, so the pharmacological expectation is that a move to another agent doing the same thing at the same transporter carries the same liability. Clinicians who have not held that distinction attribute the failure of the switch to the patient, or to the illness, rather than to the mechanism they never left.

GOOD TO REMEMBER

Because the sodium effect begins early, new confusion, unsteadiness or an unexplained fall arising in the first weeks after an SSRI is started has a pharmacological explanation sitting immediately available. Later in treatment the same presentation still deserves a sodium, but the drug is a weaker candidate for it, and reaching for the prescription first can cost the real diagnosis.

14.7 Class property or agent property: the question to ask of every finding

One question resolves most of this material under examination conditions. Given an adverse effect, a kinetic behaviour or a complaint, ask whether it belongs to serotonin reuptake inhibition or to the particular molecule doing it. The answer decides whether moving within the class can help at all, and it is the single most reliable way to reconstruct the comparison when the details have faded.

PROPERTY	CLASS OR AGENT	WHAT THE STANDARD TEXTS ANCHOR
The quantity that must be reached for effect	Class	A shared occupancy threshold rather than a shared milligram dose
Sexual dysfunction	Class	Follows serotonin reuptake inhibition itself, not the individual agent
Hyponatraemia	Class	Begins early in treatment; inappropriate antidiuretic hormone secretion offered as the probable mechanism
Muscarinic adverse effects	Agent	Paroxetine carries the greatest muscarinic affinity of the available agents
Off-target transporter and sigma binding	Agent	Sertraline adds dopamine transporter inhibition and sigma-1 receptor binding
Elimination half-life	Agent	Fluoxetine and its active metabolite are the long outliers
Discontinuation liability	Agent, read off the kinetics	More common with the shorter half-life agents

IN INDIA

Kaplan and Sadock's Comprehensive Textbook of Psychiatry, 11th edition, India chapter reports that multicentric Indian prescription-pattern studies find escitalopram and sertraline to be the most commonly prescribed antidepressants, which makes that pair the within-class comparison worth rehearsing before any other. It changes what you do with the material above. Neither of those agents is the muscarinic outlier, so a patient of yours on one of them who reports dry mouth, constipation or urinary hesitancy should send you looking for another cause rather than accepting the anticholinergic explanation that fits paroxetine. Neither is the kinetic outlier either, so the slow self-taper that fluoxetine supplies without being asked is not available for the pair you will most often be stopping, and the discontinuation question has to be raised deliberately rather than assumed away.

Sexual dysfunction and hyponatraemia belong to the mechanism; muscarinic burden, off-target binding, enzyme inhibition and half-life belong to the molecule.

15 · Beyond the SSRIs: SNRI, mirtazapine, bupropion, vortioxetine and agomelatine

Everything gathered here is defined by what it does instead of, or in addition to, blocking a serotonin transporter. The serotonin reuptake inhibitors are one pharmacological idea executed with small variations. The agents in this section are not: they share only a negative property, that none of them is a plain serotonin reuptake inhibitor, and past that point they diverge at the level of the molecule. One group inhibits two transporters rather than one. One raises monoamine transmission without touching a transporter at all. One works outside the serotonin system altogether. One adds direct receptor actions to a transporter block. One leaves the monoamine reuptake axis entirely. Examinations reward the candidate who can state, for any single agent here, the one molecular action from which the rest of its behaviour follows.

Two habits of thought are worth breaking before starting. The first is reading a class label as a description of a molecule: an abbreviation names the axis a drug was designed around, not the target it actually engages hardest, and the drugs inside one label do not engage their two transporters in the same order or to the same degree. The second is treating mechanism as a fixed property of a drug rather than of a drug at a dose, which for at least one agent in this group is simply false. Where each of these medicines belongs in the treatment of a depressive episode is the business of the Mood disorders: pharmacotherapy notes. What follows is the pharmacology that makes those decisions intelligible.

Two transporters	alpha-2	NA and DA	Melatonergic
THE SNRI AXIS, ENGAGED IN A DIFFERENT ORDER BY DIFFERENT AGENTS	MIRTAZAPINE: RELEASE RAISED, NO TRANSPORTER BLOCKED	BUPROPION, OUTSIDE THE SEROTONERGIC GROUP	AGOMELATINE, OFF THE REUPTAKE AXIS ENTIRELY

15.1 The SNRI label and the split it conceals

The abbreviation promises a drug that inhibits two monoamine transporters, and each member does, but it says nothing about which transporter the molecule reaches first or holds most tightly.

That property, the drug's **transporter preference**, is what separates the members from one another. Venlafaxine, duloxetine and desvenlafaxine all reach for the serotonin transporter first; levomilnacipran is the exception, working predominantly at the noradrenaline transporter. The label therefore conceals a genuine pharmacological split rather than a difference in potency, and a candidate who recites the class as though it were one object has already lost the discriminating detail.

The consequence is worth stating plainly. A serotonin-preferring member, at the doses in ordinary use, is doing most of its work through the same transporter an SSRI works through, with a noradrenergic component layered on top. A noradrenaline-preferring member leads with noradrenergic pharmacology and picks up serotonergic action second. Two drugs bearing one abbreviation therefore start their effect, and start their adverse effects, from opposite ends of the same axis. Which of them suits which patient is a treatment decision and belongs to the Mood disorders: pharmacotherapy notes; the pharmacological point owned here is only that the label is not the mechanism.

RULE **A class name states the axis a drug was designed around; it does not state which target the molecule engages, or at what dose it engages it.** Inside one abbreviation, one member leads with noradrenaline while the rest lead with serotonin, and at least one member changes which transporter it is meaningfully inhibiting as the dose climbs. Read the molecule, then the dose, then the label, and never in the reverse order.

15.2 Dose as a mechanism switch: venlafaxine against duloxetine

Venlafaxine is the reason mechanism has to be quoted with a dose attached. Its **dose-dependent dual action** means the serotonergic effect is present throughout the range while the noradrenergic effect is not: noradrenaline reuptake inhibition is seen only when the antidepressant dose is higher than **150 mg per day**, a threshold recorded here as mechanism because the prescribing that acts on it is owned by the Mood disorders: pharmacotherapy notes, so a prescription below that level is pharmacologically a serotonin reuptake inhibitor carrying an SNRI's name. Nothing about the drug's identity has changed; what has changed is which of its two possible actions is actually being engaged in that patient.

Duloxetine is the within-class contrast that makes the point sharp. Both transporters are inhibited from its starting doses, so the second mechanism is available without having to titrate towards it. That is a difference in kind, not a difference in potency: the two drugs do not sit at different points on one curve, they have different relationships between dose and mechanism. Stated as an examination sentence, venlafaxine acquires its second transporter with dose and duloxetine has it from the outset.

Below 150 mg per day the venlafaxine antidepressant prescription is pharmacologically a serotonin reuptake inhibitor wearing an SNRI label, a statement about mechanism and not a prescribing instruction, since prescribing is not taught here; duloxetine is dual from its starting dose.

TRAP **Answering that a patient on a low dose of venlafaxine is receiving dual-action treatment.** The error is easy to make because the class name is on the packet at every strength, and because mechanism is memorised as a property of drugs rather than of a drug at a dose. Once the threshold is crossed the dual description is accurate; below it, every line of reasoning that depends on noradrenergic action fails, including any expectation of a different adverse-effect profile from that of an SSRI. The same trap runs in reverse in a viva: asked how to obtain noradrenergic action, the answer is not to add a second drug before the first has been taken past its own threshold.

15.3 Mirtazapine: more monoamine without a transporter

Mirtazapine is the cleanest demonstration in psychiatry that raising a transmitter does not require blocking its reuptake. The mechanism begins at a presynaptic inhibitory receptor: blockade of alpha-2 noradrenergic receptors raises the release of serotonin and of noradrenaline together. Those receptors restrain transmitter output, so an antagonist there lifts a brake and release rises. Synaptic transmitter concentration climbs, as it does under a reuptake inhibitor, but by an entirely different route, and the designation the reference literature gives this drug, a **noradrenergic and specific serotonergic antidepressant**, encodes precisely that difference.

Why this distinction earns marks rather than merely being tidy: a mechanism that works by disinhibiting release behaves differently from one that works by obstructing clearance. Reuptake blockade acts wherever the transporter is expressed and depends on the transporter being there; release enhancement depends on there being tonic inhibition to lift, and therefore on the firing of the neurone itself. The two mechanisms also combine rather than duplicate, which is the pharmacological reason a release-enhancing agent and a reuptake inhibitor are not simply two versions of the same intervention. The reasoning behind combining agents in a resistant depressive illness sits with the Mood disorders: pharmacotherapy notes.

15.4 The same drug's adverse-effect signature, read off the same list

The alpha-2 blockade is only half of this molecule's pharmacology, and the other half is where the patient's complaints come from. The same compound is a postsynaptic antagonist at 5-HT2, 5-HT3 and histaminergic receptors, with a weaker action at alpha-1 adrenergic and muscarinic receptors, and the reference literature lays the drug's characteristic adverse effects at the door of that second set rather than of the presynaptic action described above. The general method of reading a receptor profile forward into a predicted adverse-effect signature is taught elsewhere in these notes; what belongs here is that this drug is its cleanest worked example, because therapeutic mechanism and adverse-effect signature sit in two separable halves of one binding profile.

- Sedation, attributed to blockade of histaminergic receptors, and the reason this agent is described as the sedating member of the group.
- Postural symptoms, from the milder antagonism at alpha-1 adrenergic receptors.
- Dry mouth and the rest of the anticholinergic complaints, from the milder muscarinic antagonism.

- The comparative absence of the adverse-effect pattern a pure serotonin reuptake inhibitor produces, attributed to blockade of 5-HT2 and 5-HT3 receptors, which prevents unopposed stimulation of those receptor populations.

Read as a whole, the profile explains a drug that raises two monoamines and yet does not behave like the drugs that raise them by transporter blockade. That is the entire teaching point, and it is reconstructible from the receptor list without memorising an adverse-effect table.

15.5 Bupropion: the antidepressant that is not serotonergic

One member of this group leaves serotonin alone. Bupropion's effects are predominantly noradrenergic and dopaminergic, and the drug is activating rather than sedating, which is the reason it is so often prescribed alongside a more sedating serotonin-enhancing antidepressant. Calling it an **activating antidepressant** is a statement about the transmitter systems engaged rather than about mood: dopaminergic and noradrenergic action drives arousal, drive and initiation, and an agent that raises those without raising serotonin has a profile unlike anything else in the antidepressant sections of this chapter. The complementary pairing follows from that fact rather than from empirical habit.

The liability that defines the molecule is seizure risk, and it is instructive precisely because it is not a fixed property of the drug. The seizure risk is dose related and is greatest with instant-release formulations, and it is lower with **slow-release** formulations at antidepressant doses **under 300 mg per day**, as stated in The Maudsley Prescribing Guidelines. Two variables therefore govern it, the amount given and the rate at which the preparation delivers it, which is a formulation-kinetics statement rather than a statement about the compound. Recognition and acute management of a drug-induced seizure belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

> **GOOD TO REMEMBER**
>
> When a seizure risk is quoted for this drug, ask which preparation the figure refers to before using it. Instant-release and slow-release preparations do not carry the same risk at the same total daily amount, so a number carried across from one formulation to the other is not conservative, it is simply wrong in whichever direction the formulations differ. The same discipline applies to any drug whose hazard is driven by peak concentration rather than by total exposure.

15.6 Vortioxetine and agomelatine: modulation past the transporter

The newest members of the group are the ones whose names carry the least information, so their mechanisms have to be learned directly. Vortioxetine is described as **multimodal**, and the word carries a precise pharmacological content: serotonin transporter blockade sits alongside antagonism at 5-HT3 and 5-HT7 and agonism at 5-HT1A and 5-HT1B. One transmitter system is therefore modulated at several points at once rather than only at its transporter. Whether that translates into a distinguishable clinical advantage is a separate question, and one this section deliberately does not answer; the pharmacology is what is owned here.

Agomelatine goes further and steps off the axis altogether. It is a **melatonin receptor agonist** and, at the same time, an antagonist at the 5-HT2C receptor, which places it outside the reuptake and the enzyme families entirely. Alone in this section, its primary target is neither a monoamine transporter nor a monoamine receptor but a melatonin receptor, and mirtazapine is the instructive contrast: that drug also raises transmission without a transporter, but it does so at a monoamine receptor, so it never leaves the monoamine system. That structural oddity is the examinable fact, and it also explains why the drug's cautions do not resemble those of its neighbours: what constrains its use is hepatic, not serotonergic. The specific liver-function testing schedule, which tests are taken, at what intervals, and the transaminase level at which the drug is stopped, is not taught here; monitoring across the psychotropics belongs to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes, and a schedule quoted from memory is exactly the kind of number that reaches a patient wrongly.

15.7 Trazodone, and dose-effect separation inside one molecule

Trazodone earns its place here for a reason that is pharmacological rather than therapeutic: it is the standard illustration that one molecule can have two dose ranges with two different purposes. Its sedative load is heavy enough that the drug is far more often written up at **50 to 100 mg**, a range the standard reference text explicitly calls subtherapeutic for antidepressant purposes, in order to treat insomnia rather than the depressive illness itself, a use whose prescribing sits with the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes. The same compound at a much higher exposure is an antidepressant; at this exposure it is a hypnotic, and calling it an antidepressant at either dose without saying which is being used is imprecise to the point of being misleading.

Generalised, the principle is that a **subtherapeutic antidepressant dose** is not a small version of a therapeutic one. It is a different pharmacological intervention, engaging whichever actions have the lowest threshold, in this case the sedating ones, and leaving the actions that require higher exposure unengaged. That is the same logic met earlier with venlafaxine, seen from the other side: there, the higher-threshold action was the one the class name advertised; here it is the one the drug is nominally licensed for.

> **PITFALL**
>
> Counting a hypnotic-dose trazodone prescription as an antidepressant trial. It appears in the medication history as an antidepressant, is often written in the evening alongside another agent, and is then carried forward into a treatment-resistance calculation as though the illness had failed to respond to it. It has not been tried at an antidepressant exposure at all, and the same reasoning error recurs with any drug whose low-dose and full-dose indications differ.

15.8 The whole set on one axis

Laid side by side, these agents answer a single question in different ways: what does the molecule do that a serotonin reuptake inhibitor does not. The accompanying comparison is best read down the middle column, because the defining action in each row predicts the entries either side of it.

AGENT	THE ACTION THAT DEFINES IT	WHAT FOLLOWS FROM THAT ACTION
Venlafaxine	Serotonin transporter first, with noradrenaline reuptake inhibition only above the threshold dose	Mechanism moves with dose, so the class name describes only the upper part of the range
Duloxetine	Both transporters inhibited from the starting dose	Dual action without titrating towards it; the within-class contrast to venlafaxine
Desvenlafaxine	Serotonin transporter preferred	Sits with venlafaxine on the serotonin-leading side of the class
Levomilnacipran	Noradrenaline transporter predominantly blocked	The one member that leads with noradrenaline rather than serotonin
Mirtazapine	Blockade of alpha-2 receptors raising release, plus postsynaptic 5-HT2, 5-HT3, histaminergic, alpha-1 and muscarinic antagonism	Monoamines rise with no transporter blocked; the postsynaptic list is the adverse-effect list
Bupropion	Predominantly noradrenergic and dopaminergic	Outside the serotonergic agents; activating rather than sedating, and paired accordingly
Vortioxetine	Serotonin transporter blockade plus 5-HT3 and 5-HT7 antagonism with 5-HT1A and 5-HT1B agonism	Modulation at several points in one transmitter system, not at the transporter alone
Agomelatine	Melatonin receptor agonism with 5-HT2C antagonism	Outside the reuptake and the enzyme families entirely; its cautions are hepatic
Trazodone	A sedative load heavy enough to separate its hypnotic use from its antidepressant use	Usually encountered well below an antidepressant exposure

MNEMONIC

Release, Activate, Multimodal, Melatonergic

The exits from the reuptake axis, in the order this section takes them. **R**elease: mirtazapine raises serotonin and noradrenaline by blocking an inhibitory adrenergic receptor. **A**ctivate: bupropion works through noradrenergic and dopaminergic effects and behaves as an activating agent. **M**ultimodal: vortioxetine adds direct receptor agonism and antagonism to a transporter block. **M**elatonergic: agomelatine leaves the monoamine reuptake families altogether.

IN INDIA

Not every agent in this section is one a trainee will actually put on a prescription. Agomelatine, vilazodone, vortioxetine and levomilnacipran are barely used in Indian practice, a pattern the Indian source literature attributes to cost or to efficacy that did not match expectation. The working consequence is where to place effort: the within-class choices that recur on Indian wards fall among the serotonin-preferring SNRIs, mirtazapine and bupropion, and those are the profiles to hold at recall depth, including the venlafaxine threshold and the bupropion formulation rule. The scarcer molecules remain fully examinable and are learned here as pharmacology rather than met at the bedside, so unfamiliarity with a drug on the ward is not a reason to leave its mechanism unlearned.

The unifying discipline, finally, is to insist on a mechanism sentence for every one of these agents before anything else is said about it. Class membership, sedation, activation and tolerability are all downstream of the molecular action, and every discriminating point an examiner can build, including the dose threshold, the transporter split, the release-versus-reuptake distinction and the formulation-dependent hazard, is recoverable from that one sentence. Cross-class comparison, interaction burden and prescribing in the medically ill are addressed in the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

16 · Tricyclics: why the class is dangerous and when it still wins

The tricyclics are the class that teaches receptor promiscuity, because almost everything that makes them dangerous is something the molecule does that has nothing to do with treating anything. A tricyclic blocks monoamine reuptake, and if that were all it did the class would look much like the selective agents that displaced it. It is not all it does. The same molecule sits on muscarinic, histaminergic and adrenergic receptors, and on a voltage-gated ion channel the heart and the brain both depend on, and those additional actions are the whole story of why a family of drugs with no efficacy problem was pushed to the back of the formulary.

The marks here are rarely for reciting a tricyclic's adverse effects. They are for the derivations underneath them: which structural feature predicts which monoamine, why prescribing one tricyclic exposes the patient to a pair of pharmacologies, which unwanted action actually kills, and why an identical dose produces very different plasma concentrations in patients of the same weight. Whether a tricyclic is the right agent for a particular patient with a depressive episode belongs to the Mood disorders: pharmacotherapy notes; overdose management, the plasma target ranges and prescribing in established cardiac disease belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

4	100 ms	40-fold	30-day
UNWANTED ACTIONS EVERY MEMBER CARRIES BEYOND REUPTAKE BLOCK	QRS WIDTH ABOVE WHICH OVERDOSE SEIZURES ARE PREDICTED	SPREAD IN PLASMA CONCENTRATION ON IDENTICAL DOSING	SUPPLY IS APPROXIMATELY THE LETHAL DOSE

16.1 The structural split that predicts which monoamine

The class divides on a single feature of the side chain, and that feature is not cosmetic: it predicts which transporter the molecule prefers. Kaplan and Sadock's Comprehensive Textbook of Psychiatry states the mapping directly. The **tertiary amines** have the greater affinity for the **serotonin transporter**, whereas the **secondary amines** primarily block the uptake of noradrenaline, so the degree of methylation on the side chain is what predicts which monoamine the drug engages.

That converts an unmanageable list into a derivation. Memorise which agent is which and the first stem naming an unrevised compound defeats you; know that the fuller methylation goes with the serotonergic profile and any compound can be placed from its structure. The split also predicts the metabolic route, so a patient's enzyme phenotype does not affect every member equally.

AXIS OF COMPARISON	TERTIARY AMINES	SECONDARY AMINES
Side-chain nitrogen	The more fully methylated form	The demethylated product of the corresponding tertiary compound
Transporter preferred	Greater affinity for the serotonin transporter	Primarily blocks the uptake of noradrenaline
Worked examples	Amitriptyline and imipramine	Nortriptyline and desipramine respectively
Standing in the body	Parent compound, converted in vivo	Also circulates as the active metabolite of its parent
Enzyme whose deficiency bites hardest	CYP2C19, which metabolises the tertiary compounds	CYP2D6, which hydroxylates desipramine and nortriptyline

16.2 Demethylation: prescribing a parent and its metabolite together

The structural split would be a tidy dichotomy if the halves of the class stayed separate in the body. They do not. **Demethylation** converts amitriptyline to nortriptyline and imipramine to desipramine, so a patient given a tertiary amine is exposed to a serotonergic parent compound and a noradrenergic **active metabolite** at the same time. The prescription names one drug; the patient receives a mixture.

Three consequences follow. The receptor pharmacology the patient experiences is the sum of parent and metabolite rather than the profile of the named compound, which is part of why the tertiary agents feel heaviest to take. The balance of that mixture is set by enzyme activity, so an inherited metabolic difference does not simply raise or lower a level, it shifts the ratio of serotonergic to noradrenergic compound. And a secondary amine is a different proposition from its parent, because a patient given the metabolite alone never receives the parent.

TRAP **Treating the tertiary and secondary amine split as a clean division of labour, so that a tertiary agent is answered as a purely serotonergic drug.** Candidates learn the mapping between structure and transporter, which is correct, and then apply it to a patient rather than to a molecule. In the patient the mapping is blurred by metabolism: the tertiary compound is being converted to the secondary one and both are circulating. A stem naming a tertiary agent and describing a noradrenergic effect is not contradicting the structural rule; it is testing whether the rule is remembered as a statement about molecules while the patient contains a mixture.

GOOD TO REMEMBER

Moving a patient from a tertiary amine to the corresponding secondary amine is not a switch to an unrelated drug. Pharmacologically it is the withdrawal of the serotonergic parent from a pair the patient was already receiving, leaving the noradrenergic compound their own metabolism had been making all along. The expectation follows: the noradrenergic action continues, the serotonergic contribution goes, and the four shared class actions stay, because the metabolite carries them too. The burden falls in size, not in kind. Whether the move is the right treatment decision for a depressive episode belongs to the Mood disorders: pharmacotherapy notes.

16.3 The unwanted actions that define the class burden

Everything so far concerns what the class is for. The rest concerns what it costs, and the cost is a set of actions contributing nothing to the therapeutic effect. Stahl's Essential Psychopharmacology puts the accusation plainly: every drug in the class shares at least four unwanted pharmacological actions in addition to reuptake inhibition, namely blockade of **muscarinic cholinergic receptors**, of H1 histamine receptors, of **alpha-1 adrenergic receptors**, and of **voltage-sensitive sodium channels**. The same source is explicit that the limitation of this class has never been its efficacy, which is the sentence that makes the topic coherent.

Separate burden from benefit that way and the adverse-effect signature stops needing to be memorised. Muscarinic blockade withdraws parasympathetic tone from the salivary glands, ciliary muscle, bladder and gut and reduces central cholinergic transmission, which is why the complaints cluster rather than arriving as unrelated symptoms. Histamine blockade blunts central arousal and disturbs appetite regulation, so sedation and weight gain travel together. Alpha-adrenergic blockade removes the vasoconstrictor response supporting blood pressure on standing, so the dizziness is postural. The general method of converting a receptor-binding profile into the complaint a patient reports is developed elsewhere in these notes; the tricyclics are simply the class in which it is least optional. The fourth action breaks the pattern, because it is not a receptor at all.

ACTION BLOCKED	PHARMACOLOGICAL CONSEQUENCE	WHAT THE PATIENT REPORTS
Muscarinic cholinergic receptors	Parasympathetic tone withdrawn from glands, ciliary muscle, bladder and gut	Dry mouth, blurred near vision, constipation, urinary hesitancy, forgetfulness
H1 histamine receptors	Central histaminergic arousal blunted, appetite regulation disturbed	Sedation, morning heaviness, increased appetite and weight gain
Alpha-1 adrenergic receptors	Vasoconstrictor support of blood pressure on standing removed	Postural light-headedness and falls in frail patients
Voltage-sensitive sodium channels	Fast inward current slowed in cardiac conducting tissue and in neurons	Little at therapeutic doses, and the lethal features in overdose

RULE **A tricyclic treats through reuptake inhibition; everything that limits the class comes from actions that are not reuptake inhibition, and one of those actions is what kills.** Efficacy and burden are carried by the same molecule through different pharmacology, which is why the later history of the antidepressants reads as an attempt to keep the reuptake block and discard the rest. Put that way, the examinable question answers itself: why a class with no efficacy problem was displaced anyway.

MNEMONIC

MASH

The unwanted actions in one word. **M**uscarinic blockade, for the dry mouth, blurred vision, constipation and confusion. **A**lpha-adrenergic blockade, for the postural drop. **S**odium channel blockade, for the conduction effect and the overdose syndrome. **H**istaminergic blockade, for sedation and weight gain. S is the odd letter: not a receptor, and the one that kills. It delivers the whole adverse-effect signature and asserts nothing beyond the source's own list.

16.4 Membrane stabilisation: the class behaves as an antiarrhythmic

The important point about the cardiac action is not that it exists but that it has a name. The Companion to Psychiatric Studies classifies it rather than describing it. Tricyclics have the properties of **Class I antiarrhythmics**, promoting **membrane stabilisation** through inhibition of fast sodium channels, and they display both Type IA quinidine-like effects, with reduced Purkinje fibre action potential amplitude, reduced membrane responsiveness and slowed conduction, and Type IB lidocaine-like effects on action potential duration and the effective refractory period.

Naming the class is not pedantry. It says this is a recognised electrophysiological action with a characterised profile, so its consequences can be predicted rather than catalogued: slowed conduction through the ventricular conducting system is what a Class I agent does, which is why the electrocardiographic change of interest here is conduction time. It also locates the pharmacodynamic danger, since adding a drug with Class I properties to another agent with Class I properties is an interaction rather than a coincidence, and the same source warns against

that combination. Interactions across drug classes, and prescribing in established cardiac disease, belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes. What belongs here is that a psychotropic sits inside a cardiological drug classification.

16.5 One channel, heart and brain, and the marker on the tracing

The sodium-channel action is present all the time, not only after an overdose, and the sources are careful about the distinction. Tricyclics block voltage-sensitive sodium channels weakly in the heart and in the brain at therapeutic doses, and in overdose that single action is held to account for coma and seizures through central nervous system effects and for arrhythmia, cardiac arrest and death through peripheral cardiac effects. One mechanism, and no separate explanation needed for the neurological features.

Weak blockade at therapeutic concentrations is tolerated because enough current remains for normal conduction and normal neuronal excitability; as concentration rises the blockade deepens in both organs, and the deepening in cardiac conducting tissue becomes measurable on the tracing before the seizures and arrhythmias it predicts.

The bedside expression of that mechanism is a measurement. Kaplan and Sadock's Comprehensive Textbook of Psychiatry gives the threshold: in overdose, a QRS interval **longer than 100 milliseconds** may predict the development of seizures and ventricular arrhythmias. The number is the mechanism made visible: the QRS complex represents ventricular depolarisation, depolarisation is the fast sodium current, and a widened QRS is a direct readout of how much channel is blocked. What follows from crossing that threshold, including the therapy it indicates, is overdose management and belongs to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

16.6 The narrow margin, and why the prescription is the hazard

The class-limiting statement in the literature is not about adverse effects at all. Goodman and Gilman's The Pharmacological Basis of Therapeutics gives the conduction effect the decisive role: the quinidine-like effect on cardiac conduction can be life threatening in overdose and is what limits use of the class in patients with heart disease, and it is the primary reason that only a limited supply should be available to the patient at any one time. Note the structure of that sentence, because it is the structure of the topic: the effect that is tolerable at therapeutic doses is the same effect that is lethal above them.

What makes that practical rather than theoretical is arithmetic. Stahl's Essential Psychopharmacology puts it bluntly: the lethal dose of a tricyclic is only about **a 30-day supply** of the drug. A routine monthly prescription is therefore, in itself, approximately a lethal quantity, which is more uncomfortable than the familiar observation that these drugs are dangerous in overdose. Nothing unusual has to happen, and that, rather than any failure of effectiveness, is why the class fell out of first-line use.

TRAP **Attributing death in tricyclic overdose to the anticholinergic burden.** The reasoning feels sound, because the anticholinergic actions dominate the clinical picture and are usually learned first. The sources place the lethal mechanism elsewhere: blockade of voltage-sensitive sodium channels is what is held to produce the coma, the seizures, the arrhythmia and the cardiac arrest, and the conduction effect is what is named as the reason the class is limited. Anticholinergic toxicity as a syndrome, and its management, belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes. The anticholinergic actions make the presentation recognisable; the sodium channel makes the overdose fatal.

16.7 Why the same dose is not the same drug

Patients of the same weight given the same tricyclic at the same dose can end up in entirely different places, and the reason is genetic rather than behavioural. Kaplan and Sadock's Comprehensive Textbook of Psychiatry sets out the pharmacogenetics for this class specifically. Roughly **5% to 10%** of Caucasians are homozygous for the recessive CYP2D6 trait and hydroxylate desipramine and nortriptyline poorly, so a standard dose behaves in them as a much larger one.

That textbook is not internally consistent about the figure, and the discrepancy is worth carrying rather than resolving. Its tricyclic chapter gives the range quoted above; its chapter on the selective serotonin reuptake inhibitors gives a narrower band, six to eight percent, for deficiency of the same enzyme in the same population. The ranges overlap and neither refutes the other, so the honest answer names this chapter's range and acknowledges the narrower one elsewhere in the same source. Other reference tables run wider still, as the cytochrome P450 section of these notes records.

The population distribution is not the same everywhere, which is where this becomes locally load-bearing. About 50% of those of Asian descent carry the **CYP2D6*10 allele**, which reduces clearance by that enzyme and makes them intermediate metabolisers. The enzyme handling the tertiary compounds is affected separately.

All of that variability arrives in the plasma, and its magnitude is the number to remember. The variability in tricyclic plasma concentration produced by these metabolic differences is substantial, and is often described as a **40-fold variation** between individuals. That is the pharmacological argument for measuring a level in this class rather than inferring one from the dose, and it is an argument about spread rather than about any particular target. The concentrations defining an adequate or a toxic level, and the framework for therapeutic drug monitoring, belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes; the enzymology itself belongs to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

IN INDIA

The pharmacogenetics of this class are not a footnote in Indian practice. About **20% of Asians** carry a polymorphism producing deficient CYP2C19 metabolism of the tertiary tricyclic compounds, and the reduced-clearance CYP2D6 allele above is carried by a large fraction of the same population. The consequence is a starting position, not a prohibition: begin at the low end, titrate against response and tolerability rather than towards a dose that suited someone else, and read a heavy anticholinergic, sedative or postural burden appearing at a modest dose as a probable metabolic phenotype rather than as intolerance to the class. One caution on the evidence: the source uses the broad category of Asian descent and reports no India-specific frequencies, so the direction is well grounded and the exact local proportion is not.

PITFALL

Choosing a tricyclic dose by analogy with the last patient who did well on one. The spread in plasma concentration between individuals on identical dosing is wide enough that a dose is a poor description of what a patient is actually receiving, and this is the class in which that gap is widest. A clinician who transfers a dose between patients, or who reads an unexpectedly heavy adverse-effect burden as the patient being difficult, is treating the dose as though it determined the exposure. The dose is an input; the response, the tolerability and where indicated the measured concentration are the output.

16.8 What the pharmacology decides, and what it does not

The class is dangerous for reasons that are entirely legible, and it survives for a reason that is equally legible. The source is explicit that what the tricyclics lost was not effectiveness; the charge against them is a set of actions bolted onto an effective molecule, not a weak molecule.

- Which monoamine an agent engages, from the methylation of its side chain, with metabolism blurring the distinction in the patient.
- Which adverse effects a patient will report, from the receptor actions accompanying the reuptake block rather than the block itself.
- Why the electrocardiogram is the relevant investigation, and why conduction is the parameter this class affects.
- Why a measured concentration adds information a prescribed dose cannot.

What the pharmacology does not settle is whether this patient should be given this drug. That judgment weighs the illness, what has already been tried and the risk the individual carries, and it belongs to the Mood disorders: pharmacotherapy notes. Pharmacology contributes one thing to it, and contributes it forcefully: in a class whose dispensed quantity is itself close to a lethal quantity, safety is not an afterthought to the prescribing decision but part of the same decision.

A tricyclic treats through reuptake inhibition and harms through everything else it does: muscarinic, histaminergic, alpha-adrenergic and voltage-sensitive sodium channel blockade, of which the sodium channel is the one that causes coma and seizures centrally and arrhythmia, cardiac arrest and death peripherally, and the lethal quantity is about a month of the prescription itself.

17 · Monoamine oxidase inhibitors: the interaction burden that defines the class

This is the only antidepressant class whose safety depends on what the patient eats. At the level of mechanism the monoamine oxidase inhibitors are the simplest antidepressants there are: they stop an enzyme from deaminating monoamines, and intrasynaptic amine concentrations rise. Everything difficult about them follows from three properties of that enzyme rather than from anything happening at a receptor. It exists in two isoenzyme forms with different substrate preferences, so which form a drug blocks decides whether it is an antidepressant at all. Most of the body's monoamine oxidase is not in the brain, and the peripheral pool works as a metabolic barrier between the gut and the systemic circulation. And when the block is covalent, function returns only when new protein has been made, which uncouples the duration of the drug's effect from the duration of the drug.

Read the class through the enzyme and its dietary rules, its waiting periods and its interaction burden all become derivable instead of memorised. What follows is that reading, from the isoenzyme split to the two structural escapes. Whether an MAOI suits a particular depressive episode belongs to the Mood disorders: pharmacotherapy notes. Recognising and treating a hypertensive crisis or a serotonin syndrome, and the cross-class interaction tables, belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

40 mg	10 mg	Up to 2 weeks	1 to 2 days
TYRAMINE DEFINING A HIGH-TYRAMINE MEAL, FASTED	FASTED TYRAMINE THAT CAN ALREADY PROVOKE A REACTION	FOR MONOAMINE OXIDASE ACTIVITY TO RETURN AFTER AN IRREVERSIBLE INHIBITOR	FOR THE SAME RECOVERY AFTER A REVERSIBLE INHIBITOR

17.1 Two isoenzymes, and what the split does and does not predict

Monoamine oxidase exists in two forms that differ in what they prefer to break down and in what inhibits them. **MAO-A** preferentially deaminates serotonin and noradrenaline and is the enzyme blocked by clorgyline and by moclobemide. **MAO-B** preferentially deaminates phenylethylamine and is the enzyme blocked by selegiline. Dopamine sits between them and is degraded about equally by both.

Two cautions attach to that tidy table and examiners like both. The first concerns noradrenaline: the Indian pharmacology standard groups it with serotonin on the MAO-A side, leaving dopamine as the only shared substrate, while the Comprehensive Textbook of Psychiatry makes dopamine and noradrenaline substrates for both forms and leaves serotonin as the only preferential MAO-A substrate. Both passages are cleanly written, so no answer should rest on the proposition that noradrenaline is an MAO-A substrate alone. The second concerns serotonin and

MAO-B. One standard pharmacology text describes MAO-B as effective against serotonin and dopamine; the psychopharmacology sources restrict it to trace amines such as phenylethylamine, with serotonin handled only at high concentrations. What survives every version is that serotonin is the secure preferential substrate for MAO-A and phenylethylamine the secure one for MAO-B.

The functional consequence is more secure than the substrate list, and it is where the teaching should be anchored. Blocking MAO-B on its own produces no antidepressant effect, because it leaves serotonin and noradrenaline metabolism untouched and, with MAO-A still working, allows almost no dopamine to accumulate. The complementary and load-bearing statement is that brain MAO-A must be substantially inhibited before an antidepressant effect occurs. That single fact organises the whole class, and it predicts in advance that any MAO-B inhibitor pushed to an antidepressant dose will have stopped being selective.

17.2 The enzyme outside the brain, and the barrier an MAOI deletes

Most teaching about antidepressants stays inside the synapse, and for this class that is the wrong place to look. **Tyramine** is produced continuously by bacterial action on dietary amino acids, both during the fermentation and ripening of food and within the gut itself. In an untreated person it never becomes a systemic problem, because monoamine oxidase in the gut wall and in the liver deaminates it before it can reach the circulation in any concentration that matters. That peripheral pool is doing nothing for mood: it is doing customs work, keeping ingested amines out of the arterial circulation.

An MAOI does not distinguish between the enzyme that raises intrasynaptic amines and the enzyme that guards the portal circulation, so the therapeutic action and the class hazard are one molecular event happening in two places. Elsewhere a drug's receptor profile predicts its adverse effects; here they are written in the enzyme's anatomical distribution instead.

RULE **The defining hazard of this class is peripheral, not central.** What an MAOI removes is the gut and hepatic barrier that normally destroys ingested amines before they reach the systemic circulation. Every characteristic danger of the class descends from that one deletion, which is why an MAOI is an interaction problem rather than a receptor problem, and why nothing in its binding profile forecasts its most dangerous events.

17.3 The pressor reaction, and the load it actually takes

The cheese reaction is a four-step chain, and setting the steps out separately shows where each of this class's rules attaches.

- Fermented foods, classically cheese, beer and red wine, are enriched in tyramine and to a lesser degree in other phenylethylamines.
- In an intact person the gut wall and the liver destroy that load on the way through, so almost none of it is absorbed systemically.
- With gastrointestinal and hepatic monoamine oxidase inhibited, a large ingested load is absorbed rapidly and reaches the systemic circulation in high concentration.

- Tyramine is not a receptor agonist but an **indirect sympathomimetic**: it is carried into sympathetic nerve terminals and displaces stored noradrenaline.

The last step is where the isoenzyme split becomes clinically decisive. Noradrenaline released by tyramine is ordinarily destroyed by MAO-A before it can do anything, and only when MAO-A is inhibited does the same release raise blood pressure. The reaction is therefore specific to MAO-A inhibition rather than to monoamine oxidase inhibition in general. When it does occur, the release of noradrenaline is massive and precipitous, and the hypertension can be severe enough to precipitate a myocardial infarction or a stroke. How that crisis is recognised and treated at the bedside belongs to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

The load required is smaller than most candidates assume. A **high-tyramine meal** is generally defined as 40 mg or more of tyramine in the fasted state, but in a patient taking an oral nonselective irreversible MAOI as little as 10 mg in the fasted state can provoke a hypertensive reaction. The fasted qualification is part of the number: the same quantity taken with a meal is better tolerated, so a threshold quoted without it is a different threshold.

17.4 Irreversibility, or waiting for protein rather than for clearance

The nonselective agents used in depression are **irreversible inhibitors**, forming a covalent complex that the enzyme cannot shed, and the pharmacokinetic consequence is the most counter-intuitive property of the class. The parent drug is excreted within 24 hours, yet monoamine oxidase activity takes up to 2 weeks to recover, because recovery depends on the synthesis and transport of new enzyme to monoaminergic nerve terminals rather than on elimination of the drug. A plasma level would therefore be useless: the patient's pharmacology is the state of the enzyme, and the enzyme has no plasma level.

Every waiting period in this class is a consequence of that arithmetic, including the one patients are least likely to observe. A low-tyramine diet has to be kept during treatment and for a further two weeks after the drug is stopped, which is exactly the time the enzyme needs to come back. The same source records that the reaction is itself rare, and rarity is why the tail of the restriction gets abandoned: nothing happened during treatment, so nothing is expected after it.

MNEMONIC

Enzyme out, enzyme in

Every interval in this class counts enzyme, never drug. The wait before a new agent is started is the time needed to make **new enzyme**; the diet that continues after stopping is that same interval seen from the patient's side. Ask when the enzyme comes back, not when the drug goes away.

TRAP **Assuming the dietary restriction ends with the last tablet.** It does not: the diet continues for a further two weeks, because what protects the patient is the enzyme, and the enzyme has not yet been replaced. Candidates err here because every other drug they know stops mattering once it is cleared, and clearance is irrelevant to this one.

17.5 The washout arithmetic, in both of its directions

Switching is where this class hurts people, and the arithmetic is not symmetrical. Coming off an MAOI, the governing quantity is enzyme resynthesis, and it is the same whatever comes next. A minimum two-week washout is required when switching from an MAOI to any other antidepressant class, and equally when switching from one MAOI to another. That second clause surprises people: when the block left behind is covalent, an MAOI-to-MAOI switch is not a cross-taper but a stop, a wait and a restart.

Where the next drug is an SSRI, the same interval is stated as **at least 14 days** after the irreversible inhibitor has been stopped, and the reason given is explicitly enzymatic: that period allows new monoamine oxidase to be synthesised. Going the other way, the governing quantity is not the enzyme but the elimination of the incoming drug and of its active metabolites, which is why the number changes with the agent. Fluoxetine is the case worth memorising: **at least five weeks** should pass between stopping it and beginning an MAOI, because Goodman and Gilman reason from a norfluoxetine half-life of one to two weeks when they set that interval. That figure is their stated premise for the washout and not a settled class number: the Comprehensive Textbook of Psychiatry gives a wider range for the same metabolite, and the comparative kinetics of the SSRIs among themselves are taken up where that class is compared.

The figure is not agreed across sources. The Maudsley Prescribing Guidelines give five to six weeks for the same switch. A shorter interval circulates in question banks drawn from a single geriatric answer key; it is the outlier, and it errs in the direction that hurts the patient. Teaching at least five weeks respects the floor that the two major sources share.

TRAP **Applying one interval in both directions.** The MAOI-first direction is governed by enzyme resynthesis and is the same for every successor drug. The MAOI-second direction is governed by elimination of the drug already on board, so it is agent-specific and, for the SSRI with the long-lived active metabolite, considerably longer. Candidates who memorise one number get one of the two directions wrong, and it is usually the dangerous one.

GOOD TO REMEMBER

Before prescribing any new psychotropic to a patient who has recently stopped an MAOI, establish the date of the last dose and count forward from it. A patient two days off an irreversible MAOI is, for prescribing purposes, still fully inhibited.

17.6 Reversibility: the escape from the diet that actually works

Moclobemide is a **reversible inhibitor of MAO-A**, and it may be prescribed without dietary precautions. Read carelessly this looks like a lesson about selectivity, and it is not: moclobemide is selective for exactly the isoenzyme whose inhibition creates the tyramine hazard, so if selectivity were protective it would be the most dangerous drug in the class rather than the safest. The protective property is reversibility. Because the inhibition is competitive, a rising tyramine load displaces the drug from the enzyme, so the substrate that would otherwise accumulate restores its own metabolism.

The washout figures prove the same point arithmetically. Full monoamine oxidase activity is restored within 1 to 2 days of stopping moclobemide, against the fortnight an irreversible agent needs, and the Maudsley switching table accordingly specifies only a twenty-four hour gap after moclobemide before an irreversible MAOI is begun. What separates those two figures is not the isoenzyme but the bond: a competitive block ends as the drug leaves, a covalent one only when new protein has been made.

How complete the protection is remains contested. The Indian pharmacology standard states flatly that potentiation of the pressor response to ingested amines is insignificant and that dietary restriction is not required. The Comprehensive Textbook of Psychiatry, writing on the same class in anxiety disorders, holds that restrictions can be less stringent but are still recommended at higher doses. Both agree on the mechanism; they differ on whether protection is complete. Treat the mechanism as settled and the practice as dose-dependent.

IN INDIA

Where an MAOI is genuinely under consideration and repeated, reliable dietary counselling cannot be assured, the reversible MAO-A inhibitor is the within-class choice that removes the problem rather than managing it, because reversible competitive block carries no dietary precaution at usual doses. That is a pharmacological argument, not a convenience one: an irreversible agent transfers a permanent responsibility to the patient's kitchen for the whole of treatment and beyond it.

Note also what the sources do not supply. The fermented-food examples in every major text are drawn from Western diets, and no source available here gives an Indian equivalent list or an Indian tyramine table. Counselling therefore has to be built from the mechanism, which is aged, fermented and bacterially ripened food of any cuisine, and not from a translated list of cheeses.

17.7 Selectivity and route: why selegiline is two different drugs

Selegiline demonstrates that isoenzyme selectivity is a property of a dose, not of a molecule. Oral selegiline at **5 to 10 mg/day** selectively inhibits MAO-B and is not effective in depression, while the antidepressant doses, not taught here, are precisely the ones at which selectivity is lost and a tyramine-restricted diet becomes necessary; the same Indian pharmacology text records the loss of enzyme selectivity at oral doses above twenty milligrams per day. Both follow from the rule that antidepressant efficacy requires substantial MAO-A inhibition: no drug is both MAO-B selective and an antidepressant.

Route of administration offers the second escape, and it works on the anatomy rather than on the chemistry. Transdermal selegiline avoids **first-pass metabolism** and is relatively sparing of the gastrointestinal MAO-A system, so brain enzyme is inhibited while the gut barrier is comparatively preserved. At the lower-strength patch, 6 mg/24 hours, tyramine restriction is not needed when the patch is used as an antidepressant, an indication owned by the Mood disorders: pharmacotherapy notes.

The higher-strength patches are contested. The Maudsley reports no hypertensive reactions even at the higher strengths, while recording that the manufacturer advises avoiding very high

tyramine foods at those strengths. Stahl's prescribing text argues that a far larger fasted tyramine load would be needed to react at the top strength, then states that studies are insufficient to be sure of safety, and elsewhere in the same book recommends food restriction at the higher transdermal doses. Where a source contradicts itself, the manufacturer caution stands.

AGENT OR ROUTE	NATURE OF THE BLOCK	ISOENZYME AFFECTED	TYRAMINE RESTRICTION	TIME TO FULL ENZYME ACTIVITY
Oral nonselective MAOI	Irreversible, covalent	Both isoenzymes	Required, and continued beyond the last dose	Up to two weeks after stopping
Moclobemide	Reversible, competitive	MAO-A	Not required at usual doses; still advised at higher doses	One to two days after stopping
Oral selegiline, low dose	Irreversible	MAO-B selective, and not antidepressant	Not applicable at this dose	As for any irreversible inhibitor
Oral selegiline, antidepressant dose	Irreversible	Selectivity lost	Required	As for any irreversible inhibitor
Transdermal selegiline, lower-strength patch	Irreversible	Selectivity lost, as at any antidepressant dose	Not required, because gut MAO-A is relatively spared	As for any irreversible inhibitor

17.8 Where the danger of this class actually sits

The dietary reaction is what the class is famous for, and it is not what most often harms people. Drug interactions are potentially more important clinically than the tyramine reactions, may be more common than the dietary interactions, and some of them can be dangerous or even lethal. That ordering should reorganise how the class is taught. The diet is visible, countable and patient-facing, so it is what gets counselled, examined and remembered; the interaction burden is invisible, sits in someone else's prescription, and is the one that kills.

Two structural features make it so. The enzyme is deleted rather than merely occupied, so the deletion persists through any drug started during those weeks; and monoamine oxidase is the disposal route for a wide range of amines, endogenous and pharmaceutical alike, so any agent that raises an amine has lost its brake. The specific pairs and the emergencies they produce belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

PITFALL

Spending the whole consultation on the food list. Clinicians counsel the diet thoroughly because it is concrete and teachable, then discharge the patient with no written instruction that nothing new, bought over the counter or prescribed by another department, is started without being checked against the MAOI. Dietary counselling is the easy half of the safety plan, and not the half with the higher body count.

Where gut MAO-A is inhibited, reversibility and not isoenzyme selectivity decides whether a patient needs a tyramine-restricted diet, and the two escapes both work by sparing that enzyme instead, one by selectivity and one by route; every waiting period in this class measures the resynthesis of enzyme, never the elimination of drug.

Mood stabilisers

18 · Lithium: the narrow-index drug and what monitoring buys

Almost every drug in this chapter is titrated against the patient; lithium is titrated against a number. That asymmetry is a pharmacological fact rather than institutional caution, and two properties of the molecule produce it. The distance between the concentration that treats and the concentration that harms is small enough that an ordinary prescribing decision, a week of vomiting or a new antihypertensive can carry a stable patient across it. And the prescription does not tell you where in that distance the patient sits, because the relation between milligrams swallowed and millimoles circulating is set by the kidney rather than by the pen. The serum concentration therefore stops being a laboratory refinement and becomes the only instrument reporting what the patient is actually exposed to. What follows develops lithium as a pharmacological object: what the ion does, how the body handles it, and what the number buys that clinical observation cannot. Lithium as a treatment for bipolar disorder belongs to the Mood disorders: pharmacotherapy notes, and the published target concentrations, with the monitoring framework built around them, belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

20 to 24 hours	28 to 48 hours	80 to 100 percent	10 to 14 hours
PERIPHERAL HALF-LIFE AT STEADY STATE	HALF-LIFE IN THE BRAIN	BIOAVAILABILITY, ANY PREPARATION	POST-DOSE SAMPLING WINDOW

18.1 Why this molecule needs a level at all

A drug earns routine blood-level monitoring when watching the patient is not a safe substitute for knowing the exposure, and lithium is the archetypal narrow-index drug because it fails that substitution twice over. Its **therapeutic index** is narrow: the gap between the concentration at which the ion works and the concentration at which it harms is too small to cross safely by guesswork. The derivation of the index, and the wider question of which psychotropics deserve a

level, are developed in the section of these notes on therapeutic index and dose-response. For lithium the index is not a comparative statistic but a constraint on every prescription.

The second failure is the more instructive. Even with a generous window you could not aim at it from the prescription, because the oral dose required to produce any given blood level varies widely between patients, and the reason it varies is renal clearance. Two adults of similar build, given the same milligrams of the same salt, sit at different concentrations, and neither history nor examination identifies which of them is high. For most psychotropics that between-patient variability is absorbed by a wide margin of safety and never has to be measured. Here the variability and the narrow margin arrive together, and the level is what closes the gap between them.

The pharmacological differences between the mood stabilisers are differences of exactly this kind: how much room the index leaves, what a level can report, and what constraint dominates prescribing. The accompanying figure sets them side by side; each agent is developed in its own section of these notes.

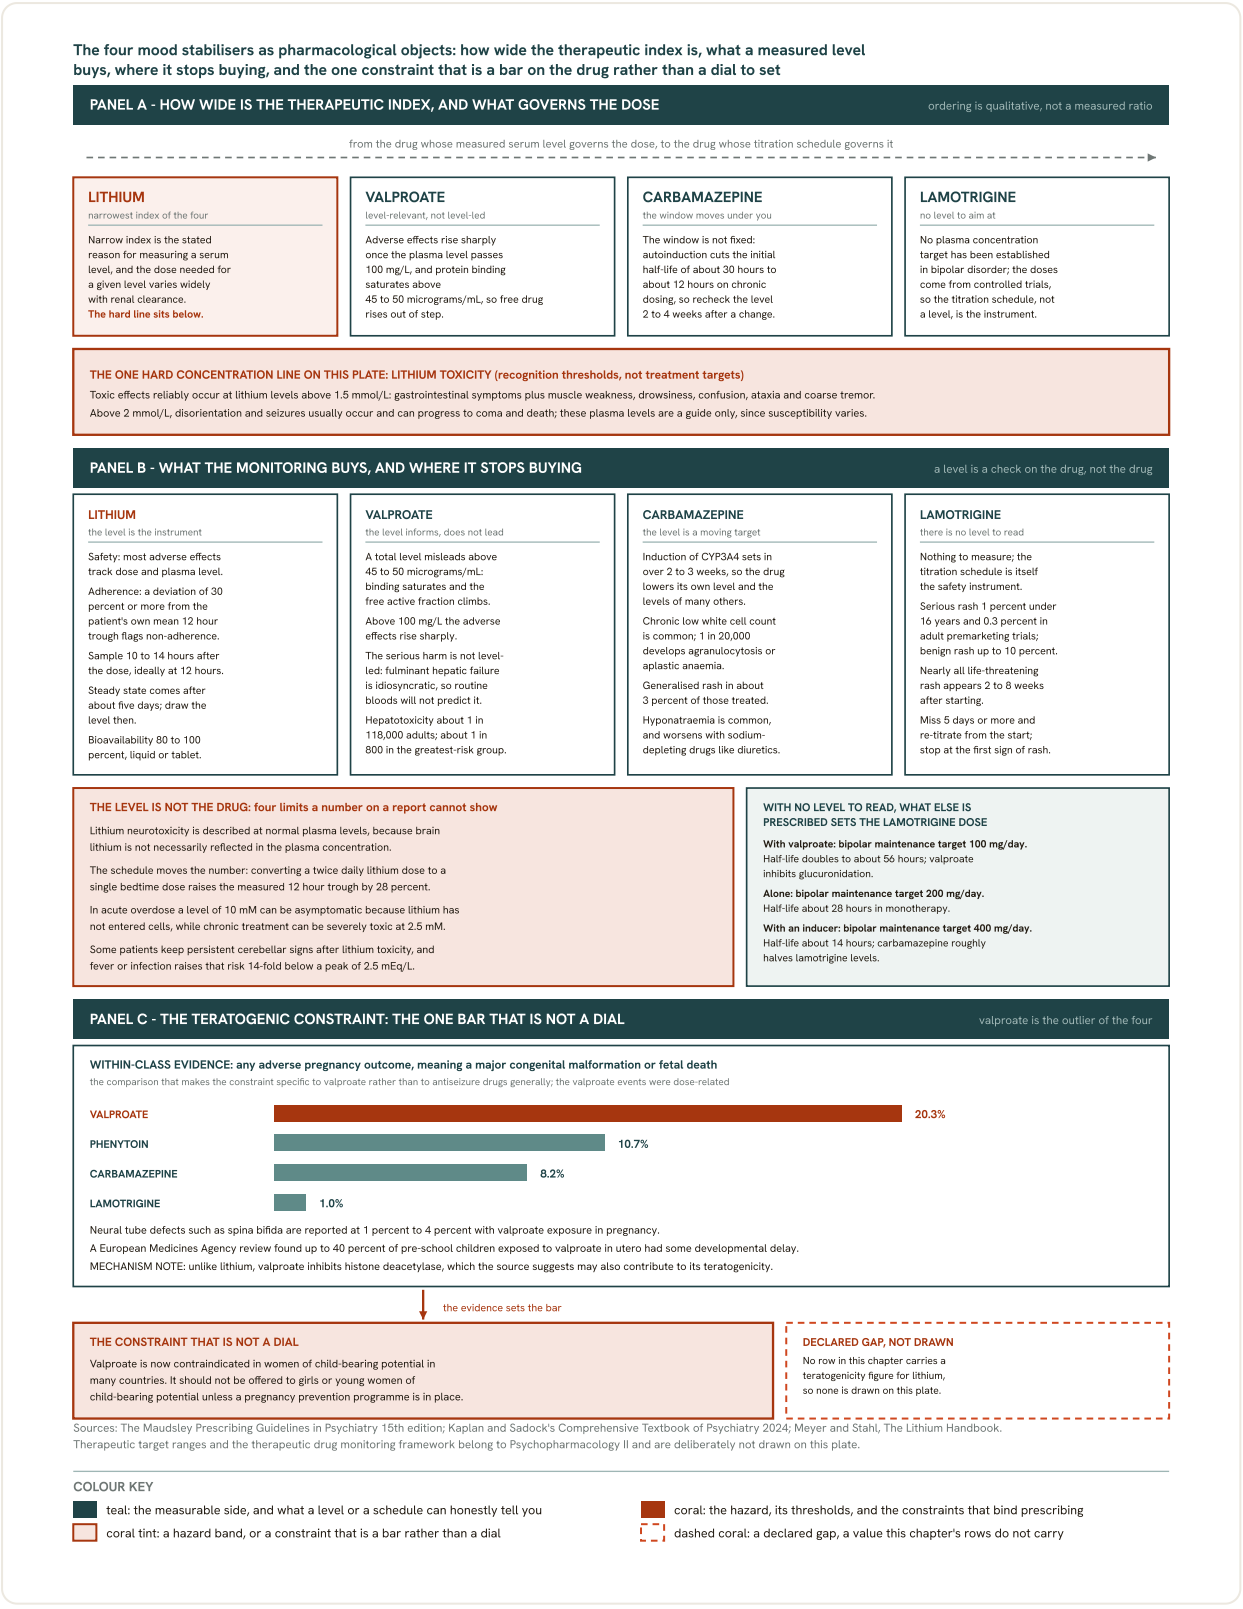


Fig 36.8 **The four mood stabilisers compared as pharmacological objects.** Lithium sits at the narrow end, where the serum level governs the dose and toxicity is reliable above 1.5 mmol/L; lamotrigine sits at the wide end, where no plasma target exists and the titration schedule is itself the safety instrument. A normal level does not exclude neurotoxicity, and the dosing schedule alone shifts the number by 28 percent. Valproate carries the one constraint that is a bar rather than a dial. (Maudsley 15th edition; Kaplan and Sadock, CTP 2024; Meyer and Stahl, The Lithium Handbook.)

RULE **A lithium dose is an instruction; a lithium level is a measurement, and only one of them describes the patient.** The dose tells you what you asked the kidney to handle; the level tells you what the kidney did with it. When they disagree, the level is right and the dose is a hypothesis.

18.2 A mechanism that has stayed plural

Lithium is the oldest drug in continuous psychiatric use and has the least settled mechanism of any agent in this chapter. That is the state of the science rather than a gap in the revision, and it is examinable as such. No single mechanism of action for lithium has been established, because the ion is implicated across so wide a range of biological processes that none has been isolated as the therapeutic route. The standard sources offer a shortlist of candidates instead, and the honest examination answer gives the shortlist and says plainly that it is one.

- **Glycogen synthase kinase 3**, an enzyme sitting at the junction of several intracellular signalling pathways.
- The cAMP response element-binding protein, which links signalling at the membrane to transcription in the nucleus.
- Pathways related to the sodium and potassium ATPase, the pump maintaining the ionic gradients lithium itself moves across.
- **Inositol monophosphatase**, the enzyme lithium inhibits inside the phosphatidyl inositol second-messenger system.

Two teaching points follow. All four candidates sit downstream of the receptor rather than at it. Every other class in this chapter is read forward from a receptor-binding profile: block this receptor, predict that adverse effect. Lithium cannot be read that way, which is why its adverse-effect signature has to be derived from distribution and renal handling rather than from a binding table. The biochemistry of the second-messenger cascades belongs to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

Second, an unsettled mechanism is not an unsettled drug: efficacy in mood stabilisation is established independently of any account of how it is produced. Candidates who present the mechanism as though one candidate had won lose marks, and so do those who present the uncertainty as though it undermined the indication.

18.3 Absorption, and what the preparations do not change

Almost nothing is lost between the tablet and the circulation. The **bioavailability** of lithium is 80 to 100 percent, and it does not depend on whether the preparation is liquid or tablet, or standard or sustained release. That is an unusually clean statement for an oral drug: switching preparations does not alter how much lithium reaches the blood over a dosing interval.

What the preparations differ in is the time taken to reach maximum concentration, which is the shape of the curve rather than the area beneath it. A sustained-release preparation flattens and delays the peak; it does not reduce the exposure. Two readings follow, pulling in opposite directions. A patient troubled by symptoms clustering shortly after the dose is a candidate for a formulation change, because the peak is what has moved. A patient whose total exposure is too

high is not helped by one, because the total exposure is what has not moved. Reaching for a sustained-release preparation to reduce exposure is a category error, and a common one.

18.4 Salt, element and ion: three numbers for one prescription

Lithium is prescribed as a salt, weighed as an element and acts as an ion, and those are three different quantities describing one tablet. The equivalence is the only route between a prescription written in one convention and a reference written in another. On the United States convention, 300 milligrams of lithium carbonate corresponds to 56 milligrams of **elemental lithium** and to 8 milliequivalents of lithium ion. United Kingdom formulations are described in millimoles instead: a 400 milligram lithium carbonate tablet carries 10.8 millimoles of lithium.

This is not pedantry. The amount of elemental lithium delivered depends on which salt the manufacturer used, so two products of nominally the same strength are not automatically interchangeable, and the Maudsley Prescribing Guidelines makes the consequence explicit: brand and strength both have to be specified, since the milligram figure alone does not identify the salt.

IN INDIA

Two decisions follow, and both are made while writing the prescription. First, name the brand and the strength rather than writing lithium and a milligram figure, because the elemental lithium delivered varies with the salt and a substitution at the counter is then a change of exposure. Second, read the unit the laboratory printed rather than the unit the textbook used: reports here follow the millimolar convention of the United Kingdom sources, while much of the reference literature is written in milliequivalents.

TRAP **Converting a lithium concentration between millimoles per litre and milliequivalents per litre.** The units genuinely differ: United Kingdom guidance reports lithium in millimoles per litre and the American texts in milliequivalents per litre or in millimolar, written mM, so a candidate who has read both has seen the same threshold written three ways and assumes a conversion factor exists. It does not. Millimolar is millimoles per litre under another name, and lithium is a monovalent ion, so all three are numerically identical and a figure quoted in one may be read directly as another. The error is not misreading the unit; it is multiplying by something to correct for it.

18.5 The kidney is the whole of the pharmacokinetics

Lithium is not metabolised in any way that matters to prescribing, so its disposition is a renal story from beginning to end. The cytochrome map that organises the interaction risk of almost every other agent here has no purchase on it, which is why an interaction with lithium is nearly always renal rather than enzymatic. The ion is freely filtered at the glomerulus, and about 70 percent of the reabsorption that follows takes place in the proximal tubule, where lithium is handled by **the sodium and hydrogen exchanger 3, or NHE3**, in direct competition with sodium.

That single sentence generates most of what a clinician has to anticipate. Because lithium and sodium compete for the same proximal transporter, the tubule cannot distinguish between them, so when sodium runs short and the proximal tubule turns avid, lithium is reabsorbed alongside it.

Sodium depletion, volume depletion and drugs acting on proximal handling therefore raise the level without a milligram being added to the prescription. A stable dose produces a stable concentration only in a stable patient, and the stability belongs to the patient rather than to the drug. The precipitants, and the recognition of an established toxic state, are developed in the section of these notes on lithium toxicity and the precipitants a prescriber controls; the renal, thyroid and parathyroid consequences of long-term exposure belong to the Endocrine and metabolic psychiatry notes.

18.6 Two half-lives, and the interval before a level means anything

Lithium has two half-lives, and confusing them produces confident errors in both directions. In adults at steady state the peripheral half-life is **20 to 24 hours**, while the half-life in the brain is longer, at 28 to 48 hours. A blood sample therefore reports the compartment that turns over quickly rather than the one in which the drug is acting, which is the kinetic basis for two observations that would otherwise look like folklore: the delay between an adequate concentration and an adequate effect, and the persistence of central effects after the serum concentration has come down.

Because the peripheral half-life is close to a full day, **steady state** is not reached quickly. Steady state arrives after five half-lives, which for lithium is about **5 days**, and a level is best drawn once it has been reached. A concentration measured earlier describes a system that has not finished responding, and it systematically understates where the patient is heading.

The central half-life also shapes how the drug should be given. Because the brain compartment turns over slowly, it is not exposed to the peaks and troughs the serum shows across a day. That long central half-life, taken together with the renal and adherence costs of dividing the dose, is the reason The Lithium Handbook gives for administering maintenance lithium as a single bedtime dose rather than in divided doses.

> **PITFALL**
>
> Drawing a level a day or two after a dose change and titrating on the result. The concentration is still climbing towards its new plateau, so the number is low, and the response to a low number is to increase the dose again. The patient then arrives at steady state on two increments instead of one, and the second answered an artefact of timing rather than a real deficit. After any change in dose, wait for the system to settle before asking it a question.

18.7 The trough is a convention, not a law of kinetics

Every published lithium concentration assumes a sampling time, and the assumption is so uniform it is rarely stated aloud. Blood for a lithium level is drawn **10 to 14 hours** after the last dose, ideally at 12 hours, in a patient taking a single bedtime dose of a prolonged-release preparation. That window is why a morning clinic and a bedtime dose fit together so conveniently, and why a sample taken at an arbitrary time cannot be compared with any published figure.

The interval has no special kinetic authority, and knowing why protects against over-interpreting it. **The twelve hour trough** became the standard because it was convenient rather than kinetically necessary, and the peripheral half-life is in fact closer to 24 hours. A drug whose half-life is around a day has not fallen far by that point, so the measured trough is not a nadir in any meaningful sense; it is a reproducible point on a slowly declining curve, and reproducibility is all a monitoring standard needs.

The corollary catches candidates and clinicians equally. If the number depends on the interval, changing the schedule changes the number without changing exposure. Converting a patient from twice daily lithium to a single bedtime dose raises the measured trough by **28 percent**, so the same total daily intake reads differently depending on how it is divided. Nothing about the exposure has moved; the sample now sits at a different point on the curve.

TRAP **Reading a rise in the trough after a schedule change as a rise in exposure.** The stem gives a patient whose total daily dose is unchanged, whose regimen has been consolidated into a single night-time dose, and whose next level is higher; the expected wrong answer is to reduce the dose. The increase is a property of when the blood was taken relative to the dose, and the correct move is to re-establish a baseline on the new schedule. The same logic condemns any comparison between two levels drawn at different intervals.

Set out together, the kinetic facts make a short table in which each line decides something.

KINETIC PROPERTY	GROUNDED FIGURE	WHAT IT DECIDES
Bioavailability	80 to 100 percent, independent of preparation	Switching preparation changes the peak, not the exposure
Proximal tubular reabsorption	About 70 percent of reabsorption	Sodium handling, not enzymes, drives interactions
Peripheral half-life	20 to 24 hours	How fast the serum compartment responds
Half-life in the brain	28 to 48 hours	Why effect lags and outlasts the serum figure
Time to steady state	Five half-lives, about 5 days	When a level after a dose change becomes meaningful
Sampling interval	10 to 14 hours post-dose, ideally 12	Whether the result can be compared with anything

18.8 What the level buys beyond safety

Monitoring is usually taught as a defensive activity, and that undersells it. Safety is the obvious purchase: most of the adverse effects of lithium are related to the dose and to the plasma level, so a patient's complaints can be interrogated as questions about exposure rather than treated as fixed properties of the drug. An adverse effect that tracks the concentration may respond to a change in the concentration, and that is a different conversation from the one ending in stopping the drug.

The second purchase is information no other instrument in psychiatry provides: an objective statement about whether the drug is being taken. Used naively that is crude, since between-patient variability means a single low value proves nothing. Used as a within-patient comparison it is informative, because a deviation of **30 percent or more** in a patient's twelve hour trough from that patient's own mean baseline level is a useful indicator of non-adherence. The comparator is the patient, not a population, which is why the first few levels on a stable regimen matter more than they appear to: they establish the baseline every later result is read against.

GOOD TO REMEMBER

Three questions make a lithium result interpretable, and asking them takes less time than the assay. When was the last dose taken, and how long before the sample was drawn. Is the patient on a single bedtime dose or a divided regimen, and has that changed since the last result. Has anything altered sodium or volume status since, including a new antihypertensive, a gastrointestinal illness or heavy sweating. A result without those three answers is a number, not a measurement.

MNEMONIC

Time since dose, Interval to steady state, Method of dosing, Estimate of adherence

Time since the last dose, which fixes where on the curve the sample sits. **I**nterval to steady state, which fixes whether the system had finished responding. **M**ethod of dosing, single or divided, which fixes what the trough means. **E**stimate of adherence, read against the patient's own baseline. All four are properties of how the sample was taken rather than of the drug.

None of this yields a target. The concentrations to aim for and the framework for acting on a result belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes. What this section establishes is why that apparatus exists: a narrow index leaving no room for approximation, a dose that does not predict a concentration, a kidney that alters the concentration without anybody altering the dose, and a sampling convention rigid enough that the number can be compared with something.

Lithium is measured because its therapeutic index is narrow and the dose needed to reach a given level varies with renal clearance; the number is interpretable only at steady state, at a fixed interval after the dose, on a known schedule, and read against that patient's own baseline.

19 · Lithium toxicity: recognition and the precipitants you control

Lithium toxicity is not an idiosyncratic reaction to the drug; it is the same drug at a higher concentration, and the concentration almost always rose because something happened to sodium. That sentence organises everything below, and it is why toxicity belongs in a systematic pharmacology chapter rather than only in an emergency one. Every other class in these notes generates its adverse-effect signature at a receptor, so the signature can be read forward from a

binding profile. Lithium offers no such profile. Its effects at high concentration are the effects of an ion that has accumulated, and accumulation is decided in the renal tubule by events that have nothing to do with psychiatry: a diarrhoeal illness, a new antihypertensive, a week of heavy sweating, an analgesic taken for a sprained ankle. The pharmacological task is therefore twofold: recognising the state and knowing what a plasma concentration can and cannot say about it, and then knowing in advance which prescriptions and which physiological changes carry a stable patient across the threshold. The management of the toxic patient, including escalation and the point at which extracorporeal removal is considered, belongs to the Mood disorders: pharmacotherapy notes. The renal and thyroid consequences of long-term exposure belong to the Endocrine and metabolic psychiatry notes, and the published target concentrations, together with the cross-class interaction matrix, belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

above 1.5 mmol/L TOXIC EFFECTS RELIABLY APPEAR	**above 2 mmol/L** SEIZURES; CAN PROGRESS TO COMA AND DEATH	**sevenfold** HOSPITALISATION RISK, ACE INHIBITOR IN THE ELDERLY	**14-fold** PERSISTENT CEREBELLAR SEQUELAE IF FEBRILE, PEAK BELOW 2.5 MEQ/L

19.1 The concentration at which the ion becomes a poison

Toxicity has a recognisable clinical shape and a concentration at which that shape becomes reliable. The Maudsley Prescribing Guidelines put it plainly: toxic effects occur dependably at concentrations **above 1.5 mmol/L**, and what appears there is a combination of gastrointestinal symptoms with central nervous system effects that includes muscle weakness, drowsiness, confusion, ataxia and **coarse tremor**. The two halves are not equally informative. The gastrointestinal component is non-specific and by itself as likely to be gastroenteritis. The neurological component makes the diagnosis, and its most useful feature is the character of the tremor: a fine tremor accompanies lithium at concentrations doing therapeutic work, whereas a coarse tremor with ataxia and slowed cognition beside it is a statement about exposure until proved otherwise.

Higher still, the picture ceases to be a matter of tolerability. **Above 2 mmol/L** disorientation increases and seizures usually occur, and the picture can progress to coma and to death. The escalation from a drowsy patient with a coarse tremor to a patient having seizures is therefore compressed into a very small distance on the concentration axis, which is what a narrow therapeutic index means in practice and why the account of lithium as the archetypal narrow-index drug sits alongside this one in these notes. The source qualifies both figures immediately, and the qualification is examinable in its own right: these plasma concentrations are a guide only, because individual susceptibility varies.

TRAP **Treating the toxicity threshold as a single number that holds for every patient.** The standard texts disagree about where toxicity begins, and the disagreement is genuine rather than an error in one of them. The Maudsley Prescribing Guidelines place the reliable appearance of toxic effects above the figure given here, which is written for the general adult patient. Geriatric Psychiatry Study Guide defines toxicity at a distinctly lower concentration and pairs that definition with a correspondingly lower target range in older patients. The two are not contradicting each other about lithium; they are describing different populations, and an older patient reaches harm at a concentration a younger one tolerates. The error is to blend them into one wide band, which yields a threshold too high for the elderly and too low for everybody else. Quote the figure with its population attached, and take the ranges themselves from the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes, which own them.

19.2 Why the number and the patient can disagree

Those thresholds are indispensable and they are also among the most over-read numbers in psychiatry, because a plasma concentration measures the wrong compartment. Lithium acts inside cells; the sample is drawn from the circulation. In equilibrium the plasma figure is an excellent proxy for exposure, and out of equilibrium it misleads in either direction. Both directions are standard examination material and they point to opposite errors.

The first is the more dangerous. Lithium neurotoxicity has been described at plasma concentrations that are entirely normal, because the concentration of lithium in the brain is not necessarily reflected in the plasma concentration. A patient who has become ataxic, dysarthric or confused on lithium therefore carries a clinical diagnosis that a reassuring laboratory report does not overturn. The report tells you what is in the serum; the signs tell you what is in the brain, and where they conflict the signs describe the compartment the drug is acting in. This is the strongest argument for examining a lithium patient rather than merely monitoring one.

The second direction runs the other way and is best seen by holding acute overdose against **chronic lithium toxicity**. In an acute ingestion a patient may arrive with a blood concentration as high as 10 mM, a concentration known to be highly toxic, and be completely asymptomatic, because the ion has not yet moved into the cells. That patient is not well; that patient is early, and the absence of signs is a timing artefact rather than reassurance. The mirror image is the patient who has taken lithium for years and drifted: severe toxicity, with diarrhoea and seizures, can appear at a much lower concentration of 2.5 mM, because the intracellular and blood concentrations are already in full equilibrium. The same laboratory value means opposite things depending on how long the ion has had to distribute, and the history rather than the assay supplies that information.

TRAP **Concluding that a normal lithium concentration excludes lithium toxicity.** The stem gives a patient on long-term lithium with new ataxia, coarse tremor or confusion, and a concentration inside the usual range, and invites the candidate to look elsewhere for a cause. The pharmacology forbids that inference, because the brain concentration is not obliged to follow the plasma concentration. The correct answer treats the clinical picture as primary and the assay as one input among several. The same logic condemns the opposite move, which is discharging an asymptomatic patient after an acute ingestion because they look well when the blood was taken.

19.3 One principle behind almost every precipitant

Learn the sodium principle and the precipitant list stops being a list to be memorised. Most of the risk factors for lithium toxicity are changes in the sodium level, or changes in the way the body handles sodium, and the categories the standard sources give fall out of that one statement:

- Low salt diets, whether prescribed for hypertension or adopted by the patient.
- Dehydration, and any state of **volume depletion**, from a diarrhoeal illness to prolonged sweating in hot weather.
- Drug interactions, which for this drug are nearly always renal events rather than enzymatic ones.
- Some uncommon physical illnesses that disturb sodium handling.

The mechanism that unifies them sits in the proximal tubule, and every precipitant below is an application of it. Lithium and sodium share the same proximal tubular exchanger, so any medication or physiological state that lowers the serum sodium causes lithium to be preferentially reabsorbed, and the concentration climbs with no change to the prescription. The tubule cannot tell the two ions apart, so when sodium becomes scarce the kidney does what it is built to do, retaining sodium avidly and retaining lithium in the same movement. That is why volume depletion from any cause deserves the same attention as a new prescription.

RULE **Almost every precipitant of lithium toxicity is a sodium event, so ask what has happened to sodium before asking what has happened to the dose.** The question sorts the whole field: a new antihypertensive, a diuretic, an anti-inflammatory, a gastrointestinal illness, a fast, a heat wave and a drug that causes hyponatraemia all arrive at the same tubular endpoint by different routes. It also tells you where to look when a concentration has risen and the prescription has not changed, which is the commonest clinical version of this problem. The answer is rarely in the drug chart column headed lithium.

19.4 Volume, and why a thiazide is not a loop diuretic

Diuretics give the cleanest demonstration, because two classes that both make a patient pass urine carry markedly different lithium risk, and the difference is entirely where in the nephron each acts. The thiazide is the dangerous one. Lithium concentrations usually rise **within 10 days** of a thiazide diuretic being started, and the magnitude of the rise is unpredictable, varying from an increase of 25 percent to 400 percent. A range that wide is not a failure of the evidence; it is the claim. No dose adjustment can be planned in advance, and the interaction has to be measured in the individual patient rather than anticipated from a table.

The loop diuretic behaves differently, and the reason is anatomical. Only 10 percent of lithium is reabsorbed in the ascending thick loop of Henle, by way of **the sodium, potassium and chloride cotransporter NKCC2**, and blocking that transporter with a loop diuretic such as furosemide presents a risk of lithium toxicity only in elderly individuals. The loop is a minor site of lithium handling, so a drug that blocks it removes a small part of the reabsorptive machinery, whereas a thiazide acts distally and provokes the compensatory proximal sodium avidity that drags lithium with it. That is the site-of-action reason behind the ranking every text gives, and it

is more useful than the ranking itself: thiazides reduce the renal clearance of lithium more than loop diuretics do.

Safer is not safe, and the ranking is about the drug rather than the patient's fluid state. Anyone rendered genuinely hypovolaemic by a diuretic is at risk whatever its class, because the precipitant is then the volume rather than the transporter.

19.5 Three prescriptions that raise a concentration without changing the dose

Beyond the diuretics, three groups account for most of the interaction risk a lithium patient carries, and none of them is a psychotropic decision. Angiotensin-converting enzyme inhibitors come first, because they combine a large effect with a slow one. ACE inhibitors produce an unpredictable rise in the lithium concentration, of up to fourfold, that develops over several weeks, and in elderly patients they raise the risk of hospitalisation for lithium toxicity sevenfold. Angiotensin II receptor antagonists may be associated with a similar risk, so switching within the renin-angiotensin blockers is no solution. The slow onset is what catches clinicians: a concentration checked promptly after the new prescription can be entirely normal and mean nothing at all.

Non-steroidal anti-inflammatory drugs come second, and are pharmacologically the most instructive of the three. NSAIDs inhibit the synthesis of **renal prostaglandins**, thereby reducing renal blood flow and possibly increasing the renal reabsorption of sodium and therefore of lithium. The effect is again unpredictable in size and in timing: reported rises range from around 10 percent to over 400 percent, with an onset that varies from a few days to several months. What makes this class clinically dominant is how it is taken. These are drugs used on a when-necessary basis and bought without a prescription, so they enter the patient's regimen without entering the drug chart, and a review that asks only about prescribed medicines will not find them.

Third are the drugs that act on lithium at one remove, by moving the serum sodium itself. Carbamazepine can cause **hyponatraemia**, which in turn leads to lithium retention and toxicity, and rare reports of central nervous system toxicity implicate the selective serotonin reuptake inhibitors for the same reason. This is the one route by which a psychotropic co-prescription becomes a lithium risk, and the mechanism is not enzymatic at all. It is recorded here as a sodium-handling event; the systematic cross-class interaction matrix belongs to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

AGENT OR CLASS	RENAL MECHANISM	EFFECT ON THE LITHIUM CONCENTRATION	TIME COURSE
ACE inhibitors, and probably angiotensin II receptor antagonists	Reduced renal perfusion with sodium retention	Unpredictable, up to fourfold	Develops over several weeks
Thiazide diuretics	Distal sodium loss with compensatory proximal reabsorption	Unpredictable, an increase of 25 percent to 400 percent	Usually within 10 days
Loop diuretics	Block NKCC2 in the ascending thick loop of Henle, where only 10 percent of lithium is reabsorbed	Smaller than a thiazide	Risk described in elderly individuals
NSAIDs	Inhibit renal prostaglandin synthesis, reducing renal blood flow	Around 10 percent to over 400 percent	A few days to several months
Carbamazepine and the SSRIs	Hyponatraemia, leading to lithium retention	Indirect, by way of the serum sodium	Variable

PITFALL

Rechecking a concentration too early after a new interacting drug and reading the normal result as an all-clear. The onset of these interactions is measured in weeks for a renin-angiotensin blocker and can run to several months for an anti-inflammatory, so a level drawn a few days after the new prescription samples a system that has not finished moving. The result is not wrong, but it answers a question nobody asked. Check when the interaction has had time to express itself, and check again the moment the patient reports neurological features rather than waiting for the next scheduled test.

19.6 What an episode of toxicity can leave behind

Toxicity is usually described as though recovery were the default and the only question were how fast the concentration can be brought down. That is optimistic. Some patients retain persistent cerebellar signs after an episode of lithium toxicity, and the picture has a name of its own, **the syndrome of irreversible lithium-effectuated neurotoxicity**. The residual signs are cerebellar rather than diffusely cognitive, so a patient who remains dysmetric or ataxic once the concentration has normalised is not simply slow to recover.

The risk factor most easily missed is not the peak concentration but the intercurrent illness beside it. Fever or infection increases the risk of persistent cerebellar sequelae **14-fold** in patients whose peak serum concentration stayed below 2.5 milliequivalents per litre. The population that figure is drawn from is the point of it: these are not the patients with the most extreme concentrations, and the multiplier applies precisely to the group whose peak concentration would not by itself have predicted lasting damage. The source states the concentration in milliequivalents per litre, which for the monovalent lithium ion is numerically

identical to millimoles per litre, as the account of lithium as a narrow-index drug elsewhere in these notes sets out. The febrile, intercurrently unwell patient who becomes toxic is therefore the one to worry about most, and the acute management of that patient belongs to the Mood disorders: pharmacotherapy notes.

19.7 The precipitants a prescriber actually controls

The sodium principle converts a long and unmemorable list into a short set of questions asked at every review. Some precipitants are physiological and can only be anticipated; others arrive on a prescription pad and are genuinely controllable. The controllable ones cluster tightly and are worth holding as a group.

MNEMONIC

ACE inhibitors, NSAIDs, Thiazides, Sodium loss

ACE inhibitors, and the angiotensin II receptor antagonists beside them. **N**SAIDs, including the ones the patient bought rather than the ones you wrote. **T**hiazides, which are the diuretic class that matters most here. **S**odium loss of any kind, whether from a low salt diet, from a diarrhoeal illness, or from a drug such as carbamazepine that lowers the serum sodium directly. ANTS suits a group of small events that arrive quietly and add up, and every letter is a change in sodium handling rather than in the lithium prescription.

GOOD TO REMEMBER

Three questions turn the pharmacology into a review that takes under a minute. First, has another prescriber added anything, with particular attention to blood pressure treatment, diuretics and analgesia, and has the patient bought anything over the counter. Second, has anything happened to fluid or salt intake: a gastrointestinal illness, fasting, a new dietary restriction, a spell of heavy sweating. Third, is the patient febrile or intercurrently unwell, which changes both the risk of toxicity and the risk that toxicity leaves something behind. A patient who answers no to all three and whose concentration has risen has an incomplete history, not a mysterious kidney.

Recognising toxicity and treating it are different tasks, and this section owns only the first: the kinetics by which an ordinary prescribing decision, made for a good reason by somebody who was not thinking about lithium, moves a therapeutic concentration into a toxic one.

Lithium toxicity is a renal event before it is a neurological one: almost every precipitant is a change in sodium or in sodium handling, the plasma concentration is a guide rather than a verdict because the brain compartment need not follow it, and the prescriptions that matter most are the ones written by somebody else.

20 · Valproate: efficacy against the teratogenic prohibition

Valproate is the one agent in this class whose pharmacology has to be learned alongside a prohibition. Nothing in its chemistry, its absorption or its tolerability at ordinary concentrations explains why an entire population of patients is now excluded from receiving it. That exclusion

rests on one body of evidence about what the molecule does to a developing fetus, and it is why its formulary entry reads as a restriction before an indication.

What follows treats valproate as a member of a drug class rather than as a treatment for anybody: mechanism, kinetics, the relation between concentration and harm, its two kinds of adverse reaction, the comparative teratogenicity that separates it from the other antiseizure mood stabilisers, and its behaviour as a metabolic inhibitor. Its efficacy in mania and in maintenance belongs to the Mood disorders: pharmacotherapy notes. Its use as an antiepileptic, and the teratogenicity question as it arises in a woman treated for epilepsy, belong to the Epilepsy and psychiatry notes. Target concentrations, the cross-class interaction matrix and the framework for prescribing in pregnancy belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

3 to 4 hours	10 to 16 hours	20.3 percent	up to 40 percent
TIME TO PEAK PLASMA CONCENTRATION	PLASMA HALF-LIFE IN ADULTS	ANY ADVERSE PREGNANCY OUTCOME	DEVELOPMENTAL DELAY AFTER EXPOSURE IN UTERO

20.1 A mechanism list, not a mechanism

Chemically this is one of the least elaborate molecules in psychiatry. Valproate is a **simple branched-chain fatty acid**, and that ordinariness is the first clue to why its pharmacology has never resolved into one account. A short fatty acid has no obvious receptor to occupy, and the sources offer instead a set of actions at several levels of the cell, none isolated as the therapeutic route.

- Inhibition of the catabolism of gamma-aminobutyric acid, raising inhibitory tone by slowing the transmitter's breakdown rather than by acting at its receptor.
- Activation of the extracellular signal-regulated kinase pathway, a signalling route linked to synaptic plasticity.
- Promotion of brain-derived neurotrophic factor expression, placing valproate among the agents credited with neurotrophic effects.
- Reduction of protein kinase C levels, a second signalling action.

Two further routes are proposed rather than established, and the distinction matters. Depletion of inositol, and indirect effects on non-GABA pathways through the inhibition of **voltage-gated sodium channels**, are put forward as further proposed mechanisms, without either being shown to be the therapeutic route. The inositol action is the interesting one for this class, because valproate shares inositol depletion with lithium: two structurally unrelated drugs converge on one second-messenger substrate. The biochemistry of those cascades belongs to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

Most of this chapter reads a drug forward from its receptor-binding profile: block this receptor, predict that adverse effect. Valproate has no binding table to read, so its adverse-effect signature

has to be derived from its plasma binding, from what its concentration does as it rises, and from a class of reactions unrelated to concentration.

RULE **Valproate produces harms of three kinds, and only two of them answer to anything measured in the patient.** Some track the concentration and improve when it falls. Some appear unpredictably in a small number of exposed patients and are anticipated by no level and no blood test. And one falls on a person who is not the patient and is fixed at the moment of exposure; the reported events were dose-related, so dose changes its size, but no dose is named at which it stops and nothing in the patient reports it. Every prescribing constraint on this drug follows from holding those three apart.

20.2 Histone deacetylase: the action lithium does not have

One action separates valproate cleanly from the other archetypal mood stabiliser, and it is the most examinable line in the mechanism. Unlike lithium, valproate inhibits **histone deacetylase**, an effect that may account for many of its actions on gene transcription. Deacetylation tightens chromatin and silences transcription; inhibiting the enzyme leaves chromatin more open and expression more permissive. A drug acting at this level is not modulating a synapse, it is modulating which genes are read.

The same source goes one step further, and the step must be reported with its uncertainty intact. It is also possible that histone deacetylase inhibition contributes to valproate's teratogenicity. That is a hypothesis offered by the source, not an established causal pathway. Its interest is structural: it is the only proposal in the mechanism list connecting how the drug works to its worst harm. A candidate who states it as fact has overstated the source; one who omits it has missed the topic's only mechanistic bridge.

20.3 Absorption, half-life and what an enzyme inducer does to them

The kinetics are straightforward, and two practical constraints follow from them. All valproate formulations are rapidly absorbed, with peak plasma concentrations reached at 3 to 4 hours, and the plasma half-life in adults is 10 to 16 hours. Rapid absorption means the formulation changes the shape of the curve rather than the total exposure. A half-life of that length means the concentration is not flat across a dosing interval, which is why several of this drug's adverse effects are described as peak-related rather than simply dose-related.

The half-life is not a fixed property of the patient. It is shortened when a **cytochrome P450 inducing drug** is co-prescribed, so a stable patient can lose exposure with no change to the prescription. Within this class that matters immediately, since the other antiseizure mood stabiliser in common use runs in the opposite metabolic direction; enzyme induction is developed in the section of these notes on carbamazepine and oxcarbazepine.

20.4 Saturable protein binding, and why the reported total misleads

This is the subtlety in valproate pharmacokinetics, and it is skipped in revision. Valproate is highly bound to albumin, and its binding sites remain unsaturated below serum concentrations of **45 to 50 micrograms per millilitre** but become saturated above that, so the pharmacologically active **unbound fraction** rises disproportionately. Above that region the albumin has no further

capacity and additional drug has nowhere to go but into the free compartment, while the laboratory continues to report bound and unbound together.

GOOD TO REMEMBER

When a valproate concentration is reported, ask what it is a total of before deciding what it means. Below the saturation region the figure tracks the active drug closely enough to serve as a proxy; above it the reported total understates how much free drug an increment has added. Increments therefore matter less than where the patient already sits, and adverse effects disproportionate to a small dose change are not necessarily exaggerated.

20.5 Two thresholds, two conventions, and neither of them a target

Two concentrations are quoted about valproate and they are easy to conflate, partly because they arrive in different units. The first is the binding threshold above; the second is a harm threshold. Many adverse effects of valproate are dose-related and often peak-level related, and they increase sharply in frequency and severity once the plasma concentration exceeds **100 milligrams per litre**. That is a statement about where tolerability deteriorates, not about where efficacy begins.

The unit problem is genuine and better stated than resolved by preference: one major source reports valproate in milligrams per litre, the other in micrograms per millilitre. Milligrams per litre is numerically identical to micrograms per millilitre, so the two thresholds read directly against each other. Side by side they say something useful: binding saturates well below the concentration at which adverse effects escalate, so a patient climbing towards the harm threshold has already spent some distance in the region where the free fraction outruns the total.

TRAP **Reading either valproate threshold as a therapeutic target range.** The stem quotes a concentration and asks what should be done next. Both figures are boundaries of quite different kinds: one marks where albumin binding saturates and the free fraction begins to rise disproportionately, the other where adverse effects escalate. Neither names a concentration to aim for. The second half of the trap is the unit convention, which differs between the sources and tempts candidates into a conversion that does not exist. Target concentrations belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

20.6 Dose-related harm and idiosyncratic harm behave differently

The most useful organising move here is to sort valproate's adverse effects by whether the concentration has anything to do with them. The dose-related group is the larger. Valproate causes dose-related tremor in up to a quarter of patients, usually an **intention or postural tremor**, a frequency high enough that it should be asked about rather than waited for. A very small proportion develop parkinsonism with cognitive decline, which reverses on discontinuation, and the reversibility is what to carry away: a parkinsonian picture in a patient on this drug is a question about the drug before it is one about a degenerative diagnosis.

The idiosyncratic group is far smaller and behaves differently. Valproate may very rarely cause **fulminant hepatic failure**, with young children receiving multiple anticonvulsants at greatest risk. The estimated risk of hepatotoxicity in adults is about 1 in 118,000, rising to about 1 in 800 in the profile the source names as greatest risk, which is polypharmacy, age under two years and

intellectual disability. Those figures come from a black-box warning table adapted from an older reference work and should be quoted as dated estimates. What survives the dating is the shape: a rare event in adults, concentrated by orders of magnitude into a young, multiply medicated, developmentally impaired group. Hyperammonaemia and pancreatitis belong to the same idiosyncratic family, rare and unpredictable, with no incidence figure quoted for either.

HARM CLASS	EXAMPLES GROUNDED HERE	RELATION TO CONCENTRATION	WHAT MONITORING CAN DO
Dose-related	Tremor, and the reversible parkinsonian picture	Dose-related and often peak-level related	A level explains the complaint and predicts the response to a reduction
Idiosyncratic	Fulminant hepatic failure, hyperammonaemia, pancreatitis	Not related to concentration	No routine test reliably anticipates the event
Developmental	Neural tube defects, developmental delay after exposure in utero	Dose-related in the reported series for malformation and fetal death; no dose relation is stated for developmental delay	Nothing measured in the patient reports the risk to the fetus

MNEMONIC

Dose-related, Idiosyncratic, Developmental

Dose-related harm answers to the concentration and can be negotiated with by lowering it. **I**diosyncratic harm does not, and is not anticipated by routine testing, so the defence is recognition rather than surveillance. **D**evelopmental harm is not a harm to the person swallowing the tablet at all, which is why no level and no liver test in the patient reports it; the reported events were dose-related, so a lower dose changes how large it is, and only non-exposure removes it.

PITFALL

Treating a schedule of liver function tests as protection against the hepatic risk. The serious hepatic event is idiosyncratic rather than dose-related, which is what makes routine monitoring an unreliable predictor of it: a normal result last month says nothing about this month, and the reassurance it generates is the problem. The same applies to hyperammonaemia, where a level checked because a protocol requires it is a different thing from one checked because a patient has become drowsy, confused or vomiting. Test in response to a change in the patient rather than to the calendar, and tell the family what change should bring them back.

20.7 The teratogenic prohibition and the comparison that grounds it

Everything above concerns the person taking the drug. The restriction that now defines valproate concerns someone else, and it rests on two harms. The first is structural: exposure in pregnancy carries an increased risk of **neural tube defects** such as spina bifida, reported at 1 percent to 4 percent. The second is neurodevelopmental. A review of studies by the European Medicines

Agency found that up to 40 percent of pre-school children exposed to valproate in utero experienced some form of developmental delay, with outcomes poorer than after carbamazepine or lamotrigine exposure. A malformation risk in the low percentages and a developmental risk reported at up to two in five of the pre-school children exposed are not the same order of statement, and it was the second that moved the regulatory position.

What makes the restriction specific to this molecule rather than to antiseizure drugs generally is the within-class comparison. The rate of any adverse outcome, meaning a major congenital malformation or fetal death, was 20.3 percent with valproate against 10.7 percent with phenytoin, 8.2 percent with carbamazepine and 1.0 percent with lamotrigine, and the valproate events were dose-related.

ANTISEIZURE AGENT	ANY ADVERSE PREGNANCY OUTCOME	POSITION WITHIN THE COMPARISON
Valproate	**20.3 percent**	Highest rate of the four
Phenytoin	10.7 percent	About half the valproate rate
Carbamazepine	8.2 percent	Lower again, an alternative in class
Lamotrigine	1.0 percent	Lowest of the four

Read that table as the argument rather than as a list. If every antiseizure drug carried a comparable risk, the response would be a class-wide caution. Because one agent sits several times above its nearest comparator and an order of magnitude above the lowest, the response is specific to that agent, and the within-class alternative is what makes a restriction workable. Valproate is now contraindicated in women of child-bearing potential in many countries, and the attached framework is explicit about the condition of any use: valproate should not be offered to girls or young women of child-bearing potential unless a **pregnancy prevention programme** is in place. Both are the United Kingdom and European regulatory position as reported by the standard prescribing guidance, and the clinical detail of operating such a programme belongs to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

IN INDIA

No position issued by an Indian regulator is asserted here, and a prescriber should establish the current local labelling rather than assume it matches the position above. What travels unchanged is the drug-level evidence, because a malformation rate, a developmental outcome and a within-class ranking are properties of the molecule and not of the jurisdiction. Three decisions follow at the point of writing the prescription, none depending on which regulator has ruled. Establish first that the indication could not be met by an agent whose comparative pregnancy outcome is lower. Second, if valproate is still the answer for a girl or a woman who could become pregnant, treat effective contraception and a recorded discussion of the teratogenic and developmental evidence as conditions of the prescription rather than good practice around it. Third, document that reasoning.

TRAP **Using the epilepsy-versus-bipolar uncertainty to soften the prohibition.** Almost all of this evidence was collected in women treated for epilepsy, and Kaplan and Sadock's Comprehensive Textbook of Psychiatry reports a cohort in which, among the small number of exposed women without an epilepsy diagnosis, the risks were not increased and were equal to those of unexposed controls; the same source notes that the risks appear more prominent in women with epilepsy than in those with bipolar disorder. Candidates read that as permission to prescribe. It is not. The Maudsley Prescribing Guidelines treats valproate as an established human teratogen and records the contraindication without reference to the indication, while warning that data drawn from epilepsy studies may not be precisely relevant to use in mental illness, which argues for caution in both directions rather than reassurance in one. Name the limitation as a limitation of how the evidence was gathered, and keep the rule.

20.8 Valproate as an inhibitor: the glucuronidation problem

The last piece of class pharmacology is what makes valproate a hazard to other drugs. Valproate raises the plasma levels of several drugs by inhibiting **glucuronidation**, lamotrigine among them. Glucuronidation is a conjugation route rather than an oxidative one, which is why this interaction is missed by clinicians who screen a new prescription only against the cytochrome map: a drug of no interest to the oxidative enzymes can still have its clearance impaired by a conjugation inhibitor.

The consequence is why the two drugs are so often examined together. A drug cleared largely by conjugation accumulates when valproate is added, and the response is a change to that drug's titration rather than to valproate's. The arithmetic for lamotrigine is developed in the section of these notes on lamotrigine and the titration that prevents the rash. The wider cross-class interaction matrix belongs to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

Set the profile out together and the shape is clear: mechanistically unresolved, kinetically well behaved apart from a saturable binding step, tolerable in most patients, dangerous unpredictably in a very small number, and hazardous to a fetus in a way no level and no liver test will detect. The contested question is not its efficacy but who may be exposed to it at all, and the teratogenic prohibition is the answer regulators reached once the within-class comparison showed the harm belonged to this agent rather than to its class.

Valproate's dose-related harms can be negotiated with by lowering the concentration and its idiosyncratic harms are anticipated by no routine test; its teratogenic and developmental harms fall on someone who is not the patient and are fixed at the moment of exposure, which is why the teratogenic prohibition restricts who may receive the drug rather than cautioning how it is given.

21 · Carbamazepine and oxcarbazepine: the enzyme-induction problem

Almost every interaction in this chapter is a transaction between two molecules; the defining interaction of carbamazepine is with itself. The drug increases the activity of the enzymes that clear it, so the concentration produced by a fixed prescription falls during the first weeks of

treatment without anybody altering the dose, and a patient adequately treated early in the course may be underexposed a few weeks later. Every other property that makes this pair examinable radiates from that single fact. The same enzyme induction that shortens the drug's own residence in the body shortens the residence of a great many co-prescribed compounds, so a stable regimen destabilises around it. And because induction is a physiological change rather than a chemical event, it takes time to develop and takes time to reverse, which means stopping the drug is as much a pharmacokinetic manoeuvre as starting it.

Oxcarbazepine belongs in the same section because it is the structural answer to that problem, and the answer is partial rather than complete. What follows develops the two agents as pharmacological objects: the channel actions they share, the potent induction only one of them has, the time course of that induction in both directions, and the adverse-effect signature that owes nothing to enzymes. Carbamazepine as a treatment choice in bipolar disorder belongs to the Mood disorders: pharmacotherapy notes, and its use as an antiepileptic to the Epilepsy and psychiatry notes. The enzymology of the cytochrome system itself is developed in the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes, while the cross-class interaction matrix, the target concentration ranges and the monitoring framework built on them belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

30 hours INITIAL PLASMA HALF-LIFE	**12 hours** HALF-LIFE ON CHRONIC DOSING	**2 to 3 weeks** TIME FOR AUTOINDUCTION TO DEVELOP	**One in 20,000** AGRANULOCYTOSIS OR APLASTIC ANAEMIA

21.1 Two molecules, one channel, and what the derivative adds

Carbamazepine acts principally on the membrane rather than at a receptor. It blocks **voltage-dependent sodium channels** and so inhibits repetitive neuronal firing, reduces glutamate release and decreases the turnover of dopamine and noradrenaline. That ordering matters for teaching: the monoamine effects are consequences of a channel action, not a separate binding profile, which is why this drug cannot be read forward from a receptor table the way an antipsychotic can. Its adverse-effect signature has to be derived from what the molecule does to excitable membranes and to hepatic enzymes, and not from occupancy at H1, M1 or D2.

Oxcarbazepine is a structural derivative of carbamazepine and keeps the channel action, but it is not simply a cleaner version of the same molecule. It works through three actions rather than one.

- Blockade of voltage-dependent sodium channels, which it shares with the parent compound.
- An increase in potassium conductance, which stabilises the membrane by a route carbamazepine does not use.
- Modulation of high-voltage activated calcium channels, altering the calcium entry that follows depolarisation.

The second structural point is about identity rather than mechanism. Oxcarbazepine acts as a pro-drug for **licarbazepine**, and eslicarbazepine, where it is available, is likewise a pro-drug for

licarbazepine. Three named agents therefore converge on one active molecule, which is worth holding because it explains why their pharmacodynamic profiles resemble one another so closely, and why a difference between them is far more likely to be kinetic than pharmacodynamic.

21.2 Autoinduction: the half-life that halves itself

Carbamazepine is a potent inducer of the cytochrome P450 oxidase enzymes, and particularly of CYP3A4. It is also a substrate of the system it induces, and those two facts together produce **autoinduction**: by increasing the amount of enzyme available, the drug increases its own clearance and reduces its own concentration. The effect is not subtle and it is not a curiosity. An initial plasma half-life of **around 30 hours** is reduced to **around 12 hours** on chronic dosing, so the elimination half-life of a patient stabilised on carbamazepine is substantially shorter than the half-life the same patient had on the first dose.

The time course is the part candidates omit, and it is the part that decides clinical behaviour. Enzyme induction requires new protein to be synthesised, so it is not immediate: autoinduction develops after 2 to 3 weeks of treatment. The consequence is a moving target. Steady state is a concept that assumes fixed clearance, and here clearance is not fixed during precisely the interval in which the clinician is trying to establish a maintenance dose. A patient can pass through an apparently satisfactory plateau early and then drift downwards as the enzyme population expands beneath them.

RULE **With carbamazepine the dose is fixed and the clearance is not, so a falling concentration on an unchanged prescription is the expected pharmacology rather than a deviation from it.** The question to ask of an unexpected carbamazepine result is therefore not what the patient did, but where in the induction time course the sample was taken, and how far the enzyme population had moved by then.

21.3 What autoinduction does to measurement

Because clearance changes during the first weeks, a concentration is only interpretable when it is anchored to a point in that time course. Plasma levels should therefore be checked **2 to 4 weeks** after starting carbamazepine or after an increase in dose, to confirm that the desired level is still being obtained. The interval is chosen to sit beyond the period in which induction develops, so that the number describes the induced state the patient will live in rather than the transient pre-induction state.

Separate two readings of that rule. A level drawn before induction is complete is not a wrong number, but it is a number about a system that has not finished changing, and it systematically overstates where the patient will settle. A level drawn after induction is complete is the one against which a maintenance dose can honestly be set. The distinction also explains why a dose increase resets the clock: the increase re-enters a period in which the enzyme population and the substrate load are adjusting to each other, and a sample taken immediately afterwards reports neither the old state nor the new one. The published concentration ranges themselves, and the wider framework of therapeutic drug monitoring, belong to the Psychopharmacology II:

emergencies, resistance, monitoring and interactions notes; what this section establishes is why a carbamazepine level has to be timed at all.

PITFALL

Reading a fall in the carbamazepine concentration during the first month as evidence of non-adherence. The clinician sees an adequate early result, a lower result a few weeks later, an unchanged prescription and an unchanged patient, and reaches for the explanation that fits almost every other drug. Here the drug itself has changed the denominator. The corrective is procedural rather than clever: record when treatment started, record when the dose last changed, and read every concentration against those two dates before reading it against the patient.

21.4 Induction as an outward-facing problem

The enzyme that carbamazepine induces does not work only on carbamazepine. Induction of the cytochrome oxidase system reduces the levels of a variety of other compounds alongside its autoinduction effect. The general principle is one every prescriber can apply without memorising a table: any co-prescribed drug cleared by the induced enzyme will fall in concentration over the same weeks, and the clinical presentation is loss of efficacy rather than toxicity. The full cross-class matrix of which psychotropics are affected belongs to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

Two consequences are properly part of this pair's pharmacology rather than of any interaction table. The first is within the mood stabiliser class itself: carbamazepine approximately halves lamotrigine levels, so a lamotrigine dose that is correct alongside carbamazepine is not the dose that would be correct alone. The titration schedule that follows from this is developed in the section of these notes on lamotrigine and the titration that prevents the rash. The second is contraceptive failure, which is a direct kinetic consequence of induction rather than a separate reproductive question: a patient established on carbamazepine who requires contraception should receive a preparation containing not less than 50 micrograms of oestrogen, or should use a non-hormonal method instead.

21.5 De-induction, or the same mechanism run backwards

Induction is a change in the amount of enzyme, so removing the inducer does not restore the previous state immediately; the enzyme population has to fall back by ordinary protein turnover. That reverse process, **de-induction**, is the half of the mechanism that examinations reward and clinics forget. When carbamazepine is stopped, for example on a switch to oxcarbazepine, de-induction of the CYP3A4 enzyme is to be expected, and substantial increases in lamotrigine levels would occur.

Notice the asymmetry in how the two directions feel to the prescriber. Starting an inducer produces falling concentrations and a slow loss of effect, which presents as a treatment that has stopped working and is easy to attribute to the illness. Stopping an inducer produces rising concentrations in every drug that was being cleared faster, which presents as new toxicity in a patient whose doses were not touched. The second is the more dangerous of the two, because nothing on the prescription changed except a drug that was taken away, and the temporal

association points in the wrong direction. Any manoeuvre involving carbamazepine, in either direction, changes the clearance of everything else the patient is taking.

TRAP **Treating the withdrawal of carbamazepine as a pharmacologically neutral act.** The stem gives a patient on carbamazepine and lamotrigine who is switched from carbamazepine to oxcarbazepine, and asks what to do about the lamotrigine. The expected wrong answer is nothing, because the lamotrigine dose has not been altered and the replacement drug looks like a near relative. The pharmacology says the lamotrigine dose was only ever correct in the presence of an inducer, that de-induction will raise its concentration substantially, and that the lamotrigine therefore needs attention at the moment the carbamazepine comes off. Candidates who have learned only that inducers lower levels answer half the question.

21.6 Oxcarbazepine: not a potent inducer, and not inert either

The derivative was designed around exactly the problem set out above, and on the central point it succeeds. Oxcarbazepine shows no autoinduction, and few major drug interactions have been reported. A clinician therefore inherits a drug whose clearance does not move underneath them, so there is no induction curve to time a level against and no drift in efficacy over the first weeks.

What the derivative does not do is escape the system altogether, and this is where the comparison earns its marks. Oxcarbazepine remains susceptible to induction by others: strong inducers of cytochrome P450 enzymes, such as carbamazepine, reduce plasma levels of its **monohydroxy derivative** metabolite by 29 percent to 40 percent. The contrast is one of degree and of direction: carbamazepine is a potent inducer and oxcarbazepine is not, and in an interaction oxcarbazepine is usually the affected drug rather than the inducing one. It has not shed inducing activity altogether, since oestrogen levels in oral contraceptives are still reduced by it.

AXIS OF COMPARISON	CARBAMAZEPINE	OXCARBAZEPINE
Principal action	Blocks voltage-dependent sodium channels, inhibiting repetitive neuronal firing	Blocks the same channels, and additionally increases potassium conductance and modulates high-voltage activated calcium channels
Active molecule	The parent compound itself	A pro-drug for licarbazepine
Effect on cytochrome P450	A potent inducer, particularly of CYP3A4	Not a potent inducer
Autoinduction	Present: an initial plasma half-life of around 30 hours falls to around 12 hours on chronic dosing	Absent
Role when inducers are involved	The inducer, of other compounds and of itself	The affected drug: strong inducers reduce the monohydroxy derivative metabolite by 29 percent to 40 percent
Hormonal contraception	For a patient needing contraception, a preparation containing not less than 50 micrograms of oestrogen, or a non-hormonal method	Oestrogen levels in oral contraceptives are still reduced

GOOD TO REMEMBER

Switching a patient from carbamazepine to oxcarbazepine solves the induction problem for the future and creates a de-induction problem for the present. Both have to be managed in the same consultation. Treat it as two separate pharmacological events: one drug leaving, which raises the concentration of everything it was clearing faster, and one drug arriving, which is stable in itself but still reduces contraceptive oestrogen and is still vulnerable to any other inducer. A switch documented as a straight substitution has recorded only half of what happened.

21.7 The rash, the allele and the ancestry question

The dermatological reaction is the other examinable axis of carbamazepine and it is entirely independent of enzyme induction. Around **3 percent** of patients treated with carbamazepine develop a **generalised erythematous rash**, and serious exfoliative reactions can rarely occur, with a vulnerability that is genetically determined. Those two statements describe different phenomena and should never be merged. The common rash is a frequent, generally benign event; the exfoliative reaction is rare, potentially catastrophic, and carries an identifiable genetic predisposition. Conflating them either terrifies clinicians out of a useful drug or, more dangerously, makes them casual about a rash that is not the common one.

The genetic predisposition is where the pharmacology becomes actionable before the first tablet. The **human leukocyte antigen (HLA) variant B*15:02** has a sensitivity of around 70 percent and a specificity approaching 100 percent in certain populations for serious skin reactions. Read

the two figures as the test's two different jobs. The very high specificity means few patients who are not susceptible test positive, so a positive result is a reason to choose another drug rather than a prediction that a reaction would have followed; with no predictive value and no population prevalence available, the post-test probability cannot be computed. The sensitivity, which is far from complete, means a negative result reduces risk without abolishing it, so the clinical vigilance that applies to every patient on this drug is not suspended by a reassuring test. On the strength of that performance, genetic testing is recommended in people from South-East Asia before carbamazepine is prescribed, a recommendation expressed by ancestry rather than by nationality or by current residence.

IN INDIA

Three decisions follow, and none of them depends on a figure this chapter does not have. Ask and record ancestry before writing the first carbamazepine prescription, because the testing recommendation issued in the United Kingdom guidance is framed by ancestry and can only be applied by a clinician who has asked the question. Where the ancestry brings the patient within that recommendation, test before prescribing rather than after a reaction. And counsel every patient started on the drug, tested or not, to stop it and present the same day if a rash appears, because a negative test still leaves a substantial minority of serious reactions undetected. The sources consulted here state no India-specific testing position and no allele frequency for Indian populations, so no such position is asserted; the ancestry recommendation is reported exactly as it is issued.

21.8 Marrow and sodium: the effects that owe nothing to enzymes

Two further adverse-effect groups complete the profile, and both are commonly misremembered because a rare event and a common one sit side by side in each. On the haematological side, carbamazepine commonly causes a **chronic low white cell count**, and one patient in 20,000 develops agranulocytosis or aplastic anaemia. These are not two grades of the same process. The benign leucopenia is frequent, usually stable and rarely a reason to stop; the marrow failure is rare and is a different event entirely. The examination distinction is the same one that separates the two rashes, and the practical distinction is that a modestly low count in a well patient is not the signal, whereas fever, mucosal ulceration or bruising in the same patient is.

On the electrolyte side, dry mouth, oedema and **hyponatraemia** are common with carbamazepine. Hyponatraemia is the one to carry forward, because it is common enough to be expected, easy to overlook when its symptoms are read as illness or as sedation, and additive with other exposures. The risk of hyponatraemia may be increased by other drugs that have the potential to deplete sodium, such as diuretics. That additivity is also what makes carbamazepine an indirect precipitant of lithium toxicity through sodium depletion; the toxic state itself, and the precipitants a prescriber controls, are developed in the section of these notes on lithium toxicity.

MNEMONIC

Self, Substrates, Stopping

Self: the drug induces its own metabolism, so its half-life falls and its own concentration drifts down over the first weeks. **S**ubstrates: the same induction lowers the concentration of other compounds cleared by the induced enzyme, which presents as loss of efficacy rather than as toxicity. **S**topping: de-induction reverses both, so withdrawal raises the concentration of everything that was being cleared faster. Oxcarbazepine deletes the first, blunts but does not abolish the second, since contraceptive oestrogen still falls, and stays susceptible to induction by others.

Carbamazepine induces CYP3A4 and is itself a substrate of it, so its initial plasma half-life of around 30 hours is reduced to around 12 hours on chronic dosing and the levels of co-prescribed substrates fall with it; oxcarbazepine shares the sodium channel action without the autoinduction, and stopping carbamazepine reverses both effects.

22 · Lamotrigine: the titration that prevents the rash

Almost every other drug in this chapter is titrated upward to find the dose that works; lamotrigine is titrated upward to avoid the dose that harms. That inversion is the organising fact of its pharmacology. The molecule is undramatic at the receptor, but the schedule by which it is introduced is a safety intervention with a named adverse outcome behind it, and it is one of the few dosing schedules a candidate must reproduce exactly. What follows develops the titration that prevents the rash as pharmacology rather than as a table to be memorised: how the body disposes of the drug, what two co-prescribed agents do to that disposal, and why every step of the schedule falls out of the kinetics.

Two boundaries are worth fixing before the detail. Lamotrigine as an antiepileptic, its titration in that setting and the rash risk as a problem of seizure treatment belong to the Epilepsy and psychiatry notes. Its place among the treatment options for bipolar disorder, including where it sits against lithium and valproate, belongs to the Mood disorders: pharmacotherapy notes. What is developed here is the systematic pharmacology both chapters assume: the disposition of the drug, the two interactions that dominate it, and the safety arithmetic that follows.

98 percent	55 percent	85 percent	10 percent
ORAL BIOAVAILABILITY	BOUND TO PLASMA PROTEIN	CONJUGATED TO INACTIVE GLUCURONIDES	EXCRETED UNCHANGED IN URINE

22.1 A mechanism that explains the effect and predicts none of the risk

The pharmacodynamic account is short. Lamotrigine **inhibits neuronal sodium channels**, and it may also inhibit glutamate release. Both are anticonvulsant actions, and the second is a plausible route to mood stabilisation given the excitatory load implicated in affective illness, but neither carries the confidence of the receptor-occupancy accounts of the antipsychotics. The

biophysics of channel gating and transmitter release belongs to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

The teaching point is what this mechanism does not do. Everywhere else in this chapter the adverse-effect signature is read forward from a binding profile: block this receptor, predict that complaint. Sodium-channel action generates no such map, and it is shared with other anticonvulsants used as mood stabilisers, so it discriminates nothing within the class. The adverse effect that governs how lamotrigine is prescribed is not a receptor effect at all: it is a rash, and no mechanistic detail about channel gating predicts it, times it or grades it.

That is why this section is shaped as it is. For an antipsychotic the receptor profile is the spine; for lamotrigine it is the disposition of the drug and the two interactions acting on it, because those set the exposure the dangerous adverse effect tracks.

22.2 A drug that is acted upon rather than one that acts

Lamotrigine is unusual among psychotropics in being almost entirely the victim of drug interactions rather than their author, and the reason is metabolic. Absorption is near complete, so the oral dose is effectively the systemic dose and no absorption variable to blame when an exposure surprises. Plasma protein binding is moderate, which keeps displacement out of the picture. The decisive fact is the route of elimination: the great majority of a dose undergoes **conjugation to inactive glucuronide metabolites**, and only a small fraction leaves the body unchanged in the urine.

Conjugation rather than oxidation is what makes this drug behave as it does in a polypharmacy regimen. The cytochrome enzymes that organise most psychotropic interaction risk are not the enzymes handling lamotrigine, which is why it is primarily a target of interactions rather than an instigator of them. What it does possess is a conjugation pathway other drugs can accelerate or block, and that asymmetry generates the rest of this section. The psychotropic substrate map is developed in the section of these notes on cytochrome P450 in practice, and the cross-class interaction matrix in the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

RULE **Lamotrigine is primarily a target of interactions rather than an instigator of them.** When a regimen containing lamotrigine misbehaves, the first question is not what lamotrigine has done to the other drug. It is what the other drug has done to lamotrigine's conjugation, and therefore to the exposure the titration was designed around. The interactions that matter here run in that one direction, and a candidate who has fixed it can reconstruct the consequences without memorising them.

22.3 One drug, three half-lives

The same tablet behaves as three different drugs depending on what else the patient is swallowing, and the variable is the elimination half-life. Taken alone, lamotrigine has a half-life of around a day and a modest clearance. Add a hepatic enzyme inducer such as carbamazepine or phenytoin and the half-life falls by about 50 percent while clearance doubles. Add valproate, which blocks the conjugation pathway, and the half-life doubles while clearance falls by

approximately 50 percent. The two interactions move the same parameter in opposite directions by comparable magnitudes, which is why neither can be treated as a minor adjustment.

Two consequences follow. The first is temporal: half-life sets how long an increment takes to declare itself, so an increase made in a patient on valproate is still climbing towards its plateau when the same increase in a patient on carbamazepine has long since settled. The second is exposure: a longer half-life at an unchanged dose means a higher steady concentration, and the steady concentration is what the serious adverse effect tracks. The general relation between half-life, steady state and the timing of a dose change is developed in the section of these notes on pharmacokinetics and half-life.

The three metabolic environments compare cleanly.

METABOLIC CONTEXT	LAMOTRIGINE HALF-LIFE	LAMOTRIGINE CLEARANCE	EFFECT ON LAMOTRIGINE CONCENTRATION
Lamotrigine alone	Around 28 hours	About 40 millilitres per minute	The reference against which the others are read
With an enzyme-inducing anticonvulsant	Approximately 14 hours	About 80 millilitres per minute	Decreased by 40 to 50 percent
With valproate	Approximately 56 hours	Around 20 millilitres per minute	Approximately doubled

22.4 Two directions, and the agents that produce them

The inducers that matter are carbamazepine, phenytoin, phenobarbital, primidone, rifampicin and oestrogen-containing hormonal contraceptives, and every one of them pulls the lamotrigine concentration down. Valproate is the single important agent pulling the other way, and it approximately doubles the concentration. The list is worth learning as a list, because the only feature its members share is an effect on the conjugation pathway.

The contraceptive entry deserves emphasis: it is the one the psychiatrist does not write. A woman established on lamotrigine who begins an oestrogen-containing contraceptive loses a substantial fraction of her exposure without any alteration to the psychiatric prescription, and one who stops it regains that fraction. Nothing has changed on the drug chart; the metabolic environment around it has. What it does not change is which column the titration is read from. The three schedules are selected by valproate and by an enzyme-inducing anticonvulsant; the prescribing information does not extend the doubled column to a contraceptive, which alters the exposure a given dose delivers and not the step the prescription opens at. Prescribing across pregnancy, lactation and the other special populations belongs to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

Enzyme induction as a class phenomenon is developed in the section of these notes on carbamazepine and oxcarbazepine, and valproate's own pharmacology in the section on valproate

as a mood stabiliser. What matters here is only the consequence for lamotrigine: the schedule is selected by the co-prescription, not by the diagnosis.

GOOD TO REMEMBER

Two questions decide which of the three lamotrigine schedules applies, and they belong before the first tablet rather than at the first review. Is the patient taking valproate? Is the patient taking an **enzyme-inducing anticonvulsant**? A third question is asked for a different reason, because it changes the exposure without changing the column: is the patient taking an oestrogen-containing contraceptive? All three are taken from the patient rather than from the referral letter, because a contraceptive is frequently absent from a psychiatric drug list. A fourth follows at every review: has any answer changed?

22.5 Two rashes, with very different frequencies and one shared instruction

The adverse effect defining this drug arrives in two forms. Benign rashes are common, occurring in as many as **10 percent** of patients taking lamotrigine. Serious rashes are rare, and they are why the prescribing information carries a **boxed warning** naming Stevens-Johnson syndrome among the reactions requiring admission and withdrawal of the drug. The incidence of serious rash was estimated at 1 percent in prospectively monitored patients under 16 years and **0.3 percent** in premarketing clinical trials of adults. That estimate dates from 1997 and should be quoted as a product-information figure rather than as a current incidence.

What matters more is the shape of the risk than its magnitude, and that shape produces an instruction that looks disproportionate until the reasoning is visible. The prescribing information does not ask the clinician to distinguish the benign rash from the dangerous one; it advises that lamotrigine ordinarily be stopped at the first sign of rash unless the rash is clearly not drug related. That rule is deliberately blunt. It accepts that most drugs stopped under it were causing nothing worse than a nuisance, because the cost of that error is a lost treatment option and the cost of the opposite error is a patient in hospital.

Nearly all life-threatening rashes have occurred between **2 and 8 weeks** after lamotrigine is started, although isolated cases have been reported after prolonged treatment. The window is the practical content of that observation: it says when vigilance must be highest, and it coincides with the period the escalation occupies. The tail is equally examinable, because a rash appearing later is not exonerated by its timing.

22.6 The risk factors a schedule can remove

Lamotrigine remains in routine use because most of the identified risk is modifiable, and the modifications are properties of how it is introduced. Besides young age, the factors raising the risk of a serious rash are increasing the lamotrigine dose too quickly, concurrent valproate, which doubles lamotrigine concentrations, and exceeding the recommended initial dose or the recommended rate of escalation.

- Young age, which the prescriber does not control and which the two incidence figures already separate out.

- Increasing the dose too quickly, which the prescriber controls completely and is the commonest way the risk is manufactured.
- Concurrent valproate, which the prescriber controls at the level of the schedule even when the co-prescription is clinically necessary.
- Exceeding the recommended starting dose or escalation rate, the same error from the other end, and the one that arrives through a prescription written from memory.

Three of the four are features of the prescription rather than of the patient, and one instrument addresses all three. That is the sense in which the titration is a safety intervention rather than a dosing convention: the deliberate removal of three named risk factors.

22.7 The schedule, and the arithmetic behind it

The prescribing information sets out three schedules, far easier to hold as one schedule with two transformations than as three lists of numbers. **The reference schedule** is the one used when the patient is taking neither valproate nor an enzyme-inducing anticonvulsant; the other two columns are read against it.

TITRATION STEP IN BIPOLAR I MAINTENANCE (PRESCRIBING OWNED BY THE MOOD DISORDERS: PHARMACOTHERAPY NOTES)	WITH VALPROATE	WITH NEITHER VALPROATE NOR AN ENZYME-INDUCING ANTICONVULSANT	WITH AN ENZYME-INDUCING ANTICONVULSANT
Opening dose (first and second weeks)	12.5 mg daily	25 mg daily, the reference step, prescribed elsewhere	50 mg daily
First increase (third and fourth weeks), indication owned elsewhere	25 mg daily	50 mg daily, a step another chapter reaches for	100 mg daily
Second increase (fifth week), not taught here	50 mg daily	100 mg daily, prescribing owned by those notes	200 mg daily
Third increase (sixth week), another chapter's decision	100 mg daily	200 mg daily, a level chosen elsewhere	300 mg daily
Target dose (seventh week onward), owned by that chapter	100 mg daily	200 mg daily, a destination not taught here	400 mg daily

The arithmetic is nearly exact, and it is the examinable part. Every step in the valproate column is one half of the corresponding step in the reference column, without exception. The enzyme-inducer column is double the reference step at every stage but the sixth week, where the reference column has already reached its target and holds while the inducer column still has one increase to

make. Nothing else changes: the step boundaries fall in the same places and the escalation takes the same number of stages in each column.

The reason is the concentration rather than the milligram. Valproate approximately doubles the concentration produced by any given dose, so halving the dose returns the patient to the exposure the reference schedule was built around. An enzyme inducer moves the concentration a comparable distance in the opposite direction, so doubling the dose does the same work. The three columns are not three levels of caution but three prescriptions calculated to deliver one exposure profile.

MNEMONIC

Halve, Double, Restart

Halve every step when valproate is in the regimen, because valproate has already doubled the exposure. **D**ouble the opening step and the target when an enzyme-inducing anticonvulsant is in the regimen, because the inducer has removed a comparable share of it. **R**estart from the opening step after any meaningful interruption, because the exposure the schedule built has gone. Three words carry the prescription once the reference column is known, with the sixth week of the inducer column read off the table rather than off the rule.

TRAP **Slowing the titration for a patient on valproate while keeping the ordinary target dose.** The stem describes a patient established on valproate in whom lamotrigine is being added, and the expected wrong answer halves the early steps and then climbs to the usual maintenance figure over a few extra weeks. It is wrong because the transformation applies to the destination as well as to the journey: the valproate column ends lower, not merely later. A doubled concentration at the ordinary target dose is the very exposure the schedule exists to prevent.

22.8 Doses that come from trials, not from a level

It is tempting to read the three columns as three attempts to hit one blood concentration, and to conclude that a lamotrigine level would settle the matter. It would not. The prescribing information derives its target doses from controlled trials rather than from a plasma concentration target, and no therapeutic range has been established for lamotrigine in bipolar disorder. The **target dose** is therefore a trial result, not a surrogate for an unstated concentration.

The consequence inverts the assumption that a level is always the more rigorous instrument. Where lithium offers a number reporting what the patient is actually exposed to, lamotrigine offers only the schedule and the dose, so everything protecting the patient is front-loaded into the prescription: the co-prescription question, the escalation rate and the response to a rash. No later measurement catches an error made at the outset. The indications for measuring a psychotropic concentration, and the published ranges for the drugs that have them, belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes; the logic of a narrow therapeutic index is developed in the section of these notes on lithium as a narrow-index drug.

22.9 An interruption resets the schedule

The last piece of the arithmetic is the one most often lost, because it is triggered by an event nobody prescribed. If a patient does not take lamotrigine for **5 days** or more, which is about five times its half-life, the drug must be reintroduced using **the full gradual escalation regimen**, because serious rashes can occur when it is restarted rapidly.

The reasoning is entirely kinetic. After roughly five half-lives a drug has effectively left the body, so a patient resuming at the maintenance dose is not continuing a titration; they are beginning one at its endpoint. Whatever the gradual introduction achieved is not a property of the prescription that can be reinstated by writing it again. The interruption is common and rarely announced.

PITFALL

Restarting a patient at the dose written on the last prescription after a gap in supply. The clinician sees a stable patient on an established maintenance dose, the pharmacy dispenses that dose, and nobody records how long the gap actually was. The question that prevents this is trivial and almost never asked: how many days did the patient go without a tablet? A gap that crosses the threshold converts a repeat prescription into a fresh initiation, and it has to be written as one.

A drug whose elimination can be doubled or halved by a co-prescription, whose dangerous adverse effect tracks exposure and appears in a defined early window, and for which no concentration can be measured, has one place left to put its safety: the schedule.

Lamotrigine's titration is a safety intervention rather than a convention: halve every step and the target with valproate, double the opening step and the target with an enzyme-inducing anticonvulsant, stop the drug at the first sign of any rash unless it is clearly not drug related, and after an interruption of several days begin the whole escalation again.

Anxiolytics, hypnotics and other agents

23 · Benzodiazepines: tolerance, dependence and withdrawal as pharmacology

Almost everything clinicians disagree about with the benzodiazepines is settled at the receptor. Why the class is so forgiving taken alone and so dangerous taken in company, why one member sedates while another mainly relaxes muscle, and why the last part of a reduction is harder than the first, are all consequences of one receptor complex and of what a benzodiazepine molecule does to it. This topic sets out that pharmacology and then reads tolerance, dependence and withdrawal off it, as receptor-level phenomena with a mechanism rather than as failures of prescribing or of character.

The ground is examined heavily because it is one of the few places in psychopharmacology where a mechanism, a within-class distinction and a well-quantified clinical consequence sit in a straight line, so one question can test all three. The use of these drugs as treatments for anxiety and for

insomnia belongs to the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes; benzodiazepine dependence as a diagnosable substance use disorder, with its assessment and its detoxification, belongs to the Substance use disorders notes; and the biochemistry of GABA as a transmitter belongs to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes. What is owned here is the class as a pharmacological object.

58 to 100%	15 to 44%	up to 15%	alpha-1
OF LONG-TERM USERS HAVE A WITHDRAWAL REACTION	HAVE MODERATE TO SEVERE WITHDRAWAL ON STOPPING	DEVELOP A PROTRACTED WITHDRAWAL SYNDROME	THE SUBUNIT CARRYING SEDATION AND ATAXIA

23.1 The receptor, and what the drug actually does to it

The target is the GABA-A receptor, a ligand-gated chloride channel. It is a **pentamer**, and the assembly most commonly found in the brain is built from **two alpha subunits, two beta subunits and one gamma subunit**, which is the common arrangement rather than the only one. Benzodiazepines do not bind where GABA binds. They bind at **the interface of the gamma-2 and alpha subunits**, a site of their own, and that structural separation is the reason the class can only ever modify a signal that is already arriving.

The structural point becomes the defining pharmacological one. A benzodiazepine binds directly to a specific modulatory site on the receptor and enhances the effect of GABA, but it does not open the chloride channel in the absence of GABA. It is therefore a **positive allosteric modulator** and not a direct agonist. That is not hair-splitting, it is the origin of the class's safety profile: the drug supplies no inhibition of its own, it only amplifies inhibition the brain is already generating. What the allosteric change does is increase the **frequency with which the central chloride channel opens**. More openings per unit time, on a channel that will not open at all unless the transmitter is present.

The barbiturates are the instructive contrast, and it is architectural rather than a matter of strength. They act at a different site on the same complex, and at higher concentrations they become direct GABA-A agonists. A drug that can open the channel unaided has no ceiling written into its mechanism; a drug that can only amplify an existing signal has one built in before the first prescription. Two groups producing a similar gross clinical effect are therefore not pharmacologically interchangeable, and the question is settled at the binding site.

RULE **A benzodiazepine modulates; it does not open.** Every distinctive property of the class descends from that one sentence. The safety in isolated overdose follows from it, because the effect is bounded by how much GABA the brain is releasing. So does the shape of withdrawal, because a system that has adapted to chronic amplification of its own inhibition is left under-inhibited when the amplifier is removed.

TRAP **Writing that a benzodiazepine increases the duration of chloride channel opening.** The action is on the *frequency* of opening, and frequency is a distinct parameter from duration. Candidates memorise the benzodiazepine and barbiturate pair as a matched contrast, then reconstruct it under time pressure from whichever half they recall first. Learn the benzodiazepine half on its own terms, as frequency of opening at a modulatory site that requires GABA to be present, and the answer stops depending on the other half.

23.2 Subunit selectivity, and why one class has several personalities

Not every GABA-A receptor responds to a benzodiazepine, and the discriminating component is the alpha subunit, of which there are six numbered subtypes. Receptors carrying alpha-1, alpha-2, alpha-3 or alpha-5 are **benzodiazepine-sensitive**, while the presence of **alpha-4 or alpha-6** leads to a lack of sensitivity to the drugs altogether. A fraction of the inhibitory machinery is therefore invisible to the class, worth holding whenever an effect is attributed to global GABA enhancement.

Among the sensitive receptors, the subunit predicts which effect is produced. Receptors containing alpha-1 appear to mediate the sedative and hypnotic effect, and the same subunit is implicated in ataxia. The antianxiety effect is attributed instead to receptors containing **alpha-2 and possibly alpha-3**, and those two subunits also carry the muscle relaxant effect. The hedging is the source's own and should be preserved: the alpha-3 attribution is stated as a possibility, not a finding.

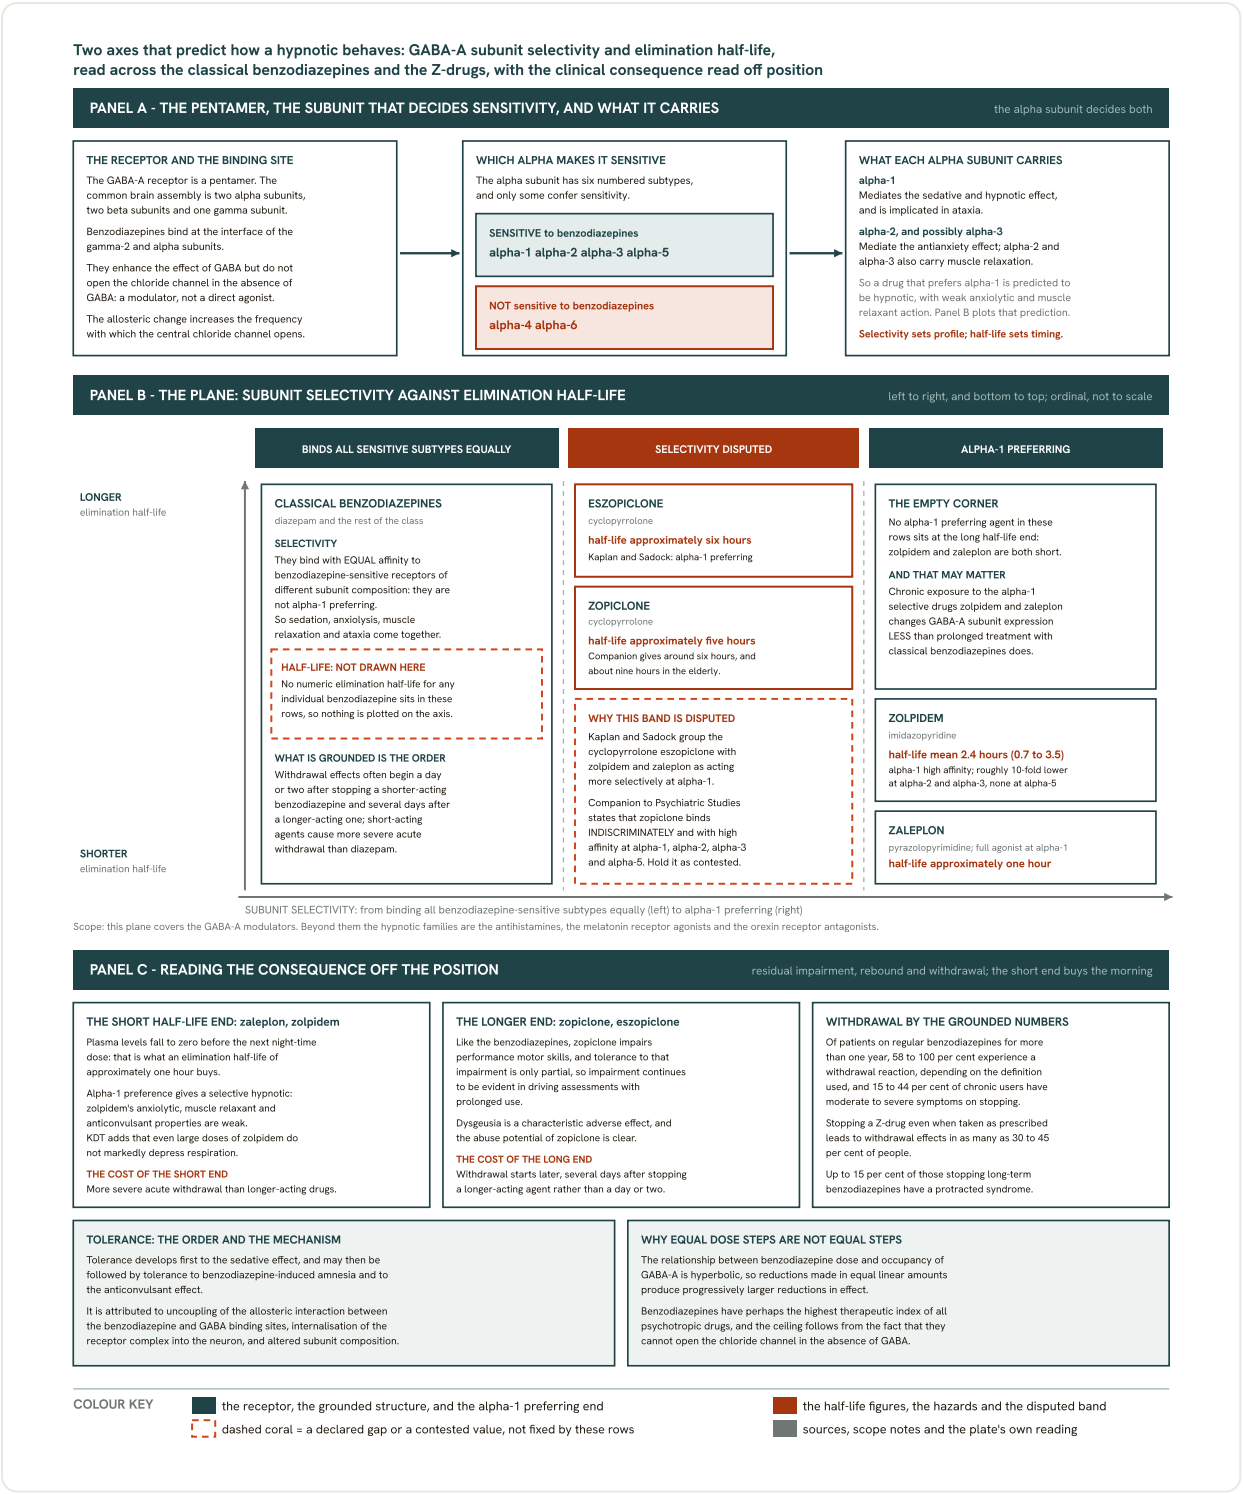


Fig 36.9 GABA-A subunit selectivity set against elimination half-life. Subunit preference predicts the profile and half-life predicts the timing: alpha-1 carries sedation, hypnosis and ataxia, while alpha-2 and possibly alpha-3 carry anxiolysis and muscle relaxation. The classical benzodiazepines bind benzodiazepine-sensitive receptors of different subunit composition with equal affinity and are drawn without a half-life, because no numeric half-life for an individual benzodiazepine sits in the rows behind this plate; the Z-drugs are plotted in half-life order, with the cyclopyrrolones marked as the contested selectivity case. The short end buys a clear morning and pays in more severe acute withdrawal, the longer end carries residual motor impairment and a later withdrawal onset. (Kaplan and Sadock, CTP 2024; Companion to Psychiatric Studies; The Maudsley Deprescribing Guidelines; KDT Essentials of Medical Pharmacology.)

ALPHA SUBUNIT PRESENT	BENZODIAZEPINE SENSITIVITY	EFFECT ATTRIBUTED TO IT
alpha-1	Sensitive	Sedative and hypnotic effect; implicated in ataxia
alpha-2	Sensitive	Antianxiety effect, and muscle relaxation
alpha-3	Sensitive	Possibly antianxiety; also muscle relaxation
alpha-5	Sensitive	No effect attributed to it in these sources
alpha-4 and alpha-6	Not sensitive	No benzodiazepine action at these receptors

Two consequences follow, and the accompanying figure lays the same logic out alongside the kinetic axis. The first is pharmacological ambition: if sedation and anxiolysis are carried by different subunits, a subunit-selective molecule could in principle separate them, which is the reasoning behind every attempt to build an anxiolytic that does not sedate. The second is diagnostic. When one patient is unsteady and another on a different member is chiefly calm, that is not idiosyncrasy, it is a difference in where the molecule works laid over a difference in how long it stays there. Half-life and steady state have their own section in these notes, as does the comparison with the Z-drugs, which exploit precisely this subunit axis.

23.3 The therapeutic index, and the one route to respiratory harm

The mechanical ceiling has a measurable consequence. Benzodiazepines have perhaps the highest **therapeutic index** of all psychotropic drugs, a remarkable position for a class this sedating to occupy. It is the arithmetic of a modulator: escalating the dose recruits more amplification of an inhibitory signal that is itself finite, so the curve flattens where an agonist's would keep climbing. What a narrow index means in general, and which psychotropics consequently need a concentration measured, has its own section in these notes.

The ceiling does not abolish respiratory risk; it relocates it. At therapeutic doses the class does not interfere with respiration, but rapid intravenous administration of a swiftly distributed compound such as diazepam can produce significant respiratory depression, and this is reversible with the benzodiazepine antagonist **flumazenil**. Route and speed, not dose alone, generate the hazard, and because the drug occupies a defined site another molecule can displace it from that site, a reversibility no metabolic strategy could deliver.

The historical comparison finishes the argument. Barbiturates are blighted by high dependency potential and a very narrow therapeutic index, and the same problems afflicted their analogues glutethimide and methaqualone, popular in the 1960s and 1970s. The benzodiazepines displaced them not by being better sedatives but by having a mechanism with a ceiling, and the property that made the class safe is also what made it prescribable at the scale that created the dependence problem.

GOOD TO REMEMBER

The high therapeutic index is a statement about the drug alone. It does not survive combination with other central nervous system depressants, because the ceiling is on this drug's modulation of GABA, not on total central depression from every source acting at once. A patient told the tablets are very safe has been told something true about a situation that does not describe them if they also drink, or if another prescriber adds a second sedating agent. The sentence that carries the pharmacology accurately is that the drug is safe by itself, conditional on staying by itself. Interactions across drug classes are developed in the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

23.4 Tolerance: the order it arrives in, and what moves at the receptor

Tolerance to this class is not a single event but a sequence, and the sequence is examinable. It develops to many of the drugs' actions but not all, and it develops first to **the sedative effect**. That may then be followed by tolerance to benzodiazepine-induced amnesia and to the anticonvulsant effect. The clinical reading is immediate: the effect a patient most notices at the start is the effect most likely to disappear, which is the pattern that produces dose escalation in someone who is not drug-seeking but simply chasing back an effect that has faded.

MNEMONIC

Sedation first, then amnesia, then anticonvulsant

Sedation goes first and goes fastest, which is why the initial sleepiness is a poor guide to how the drug will feel later. **A**mnesia may follow. **A**nticonvulsant effect may follow as well, and that is the limb with the sharpest consequence, because an agent relied on for seizure control cannot be assumed to hold that effect indefinitely. The order runs from the effect the patient reports to the effect nobody notices until it is needed.

Underneath the sequence there is a receptor-level account, and the honest version is that the sources call the mechanism unclear and describe it as linked to dose and duration of exposure. What is named is a set of adaptations:

- **Uncoupling** of the allosteric interaction between the benzodiazepine and GABA binding sites, so the drug's occupancy stops translating into enhancement.
- Internalisation of the GABA-A receptor complex into the neuron, which removes receptor from the membrane altogether.
- Alteration in the subunit composition of the receptors, which changes what the surviving population responds to.

Learn them together: they explain why tolerance is not reversed by a dose increase, and why the receptor a patient has after long exposure is not the one they started with. Receptor up-regulation, down-regulation and sensitisation as general phenomena belong to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

23.5 Withdrawal is the ordinary case, not the exception

The single most common misconception about this class is that withdrawal is what happens to a minority who misuse it. The figures say otherwise, and they are worth carrying away. In patients who have taken regular benzodiazepines for more than one year, 58 to 100 per cent experience a withdrawal reaction, the width of that band reflecting how differently studies define the event, not any doubt that it occurs. Narrowing to clinically substantial symptoms, **15 to 44 per cent** of chronic users experience moderate to severe withdrawal symptoms on stopping.

Read those two numbers as one statement with two thresholds. Some degree of withdrawal after a year of regular exposure is the expected pharmacological outcome, while a substantial minority have symptoms severe enough to dominate the picture. Neither figure describes a substance use disorder. They describe physiological adaptation in patients taking the drug as prescribed, and conflating the two attributes to the patient's conduct what belongs to the receptor. The disorder, its diagnosis and its treatment belong to the Substance use disorders notes.

> **PITFALL**
>
> Reading withdrawal as relapse. When anxiety, insomnia and autonomic arousal appear during or after a reduction, they overlap almost entirely with the presentation the drug was started for, so the natural inference is that the underlying disorder has returned and the dose should go back up. Given how commonly a withdrawal reaction follows long-term use, the prior probability sits the other way. The discriminators are temporal and pharmacological rather than phenomenological: the relationship of onset to the reduction, and the half-life of the agent reduced. Getting this wrong does not merely delay a taper, it teaches the patient that they cannot manage without the drug.

23.6 The timetable is written by half-life, but not the whole of it

Onset of withdrawal tracks elimination, the most directly useful kinetic prediction here:

- After a shorter-acting benzodiazepine, effects often begin a day or two after stopping.
- After a longer-acting one they often begin several days after stopping, and delayed-onset effects are also recognised.

Severity follows the same axis in the acute phase. Short-acting benzodiazepines are associated with more severe acute withdrawal than longer-acting drugs such as diazepam, which is what the pharmacology predicts: a concentration falling steeply gives the adapted system no time to readjust, while a slow decline is a taper the drug performs on the patient's behalf. That principle explains most of the within-class variation clinicians encounter, and it is the content behind every equivalence table, since potency and half-life vary independently and a table is needed precisely because knowing one tells you little about the other. The numeric equivalences are approximations rather than conversions, and belong to the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes.

What half-life does not predict is the long tail. **A protracted withdrawal syndrome** is estimated to occur in up to 15 per cent of patients stopping long-term benzodiazepines, lasting months and

sometimes years, and both short and long acting agents can produce it. Regulation has followed the evidence: in 2020 the FDA boxed warning on marketed benzodiazepines was updated to state that withdrawal syndromes could persist for **more than 12 months** after stopping.

TRAP **Concluding that a long half-life protects against withdrawal.** It softens the acute phase, which is real and useful, and it does nothing for the protracted syndrome, which both short and long acting agents can produce. Candidates who compress the two findings into a single rank order answer confidently and wrongly. Hold them apart as two questions: how severe the acute phase is, which half-life governs, and whether a syndrome persists for many months, which half-life does not govern.

23.7 Why equal dose steps are not equal effect steps

The last piece of mechanism is the one most often missing from a candidate's account, and it explains why the end of a reduction behaves differently from the start. The relationship between benzodiazepine dose and occupancy of GABA-A is **hyperbolic**, not linear. Occupancy rises steeply at low doses and then flattens, so most of the receptor effect has already been bought by a small fraction of a typical dose.

Run the curve backwards and the clinical implication appears. Reductions made in equal linear amounts produce progressively larger reductions in effect, because each identical decrement moves further down the steep part of the curve than the one before it. A patient who tolerated the early reductions and then struggles badly on the last ones is not deteriorating or failing; they are meeting the geometry of the curve, which the pharmacology predicted before the reduction began. The corollary is that a plan reasoned in equal steps of milligrams is reasoning in the wrong currency, because the currency that matters is receptor occupancy. The schedules that follow belong to the prescribing and deprescribing lens rather than to class pharmacology, and are found in the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes.

23.8 The edges of what class pharmacology settles

Two limits deserve stating plainly, because both invite a manufactured answer. The first is that this account is built on classical benzodiazepines acting at benzodiazepine-sensitive receptors, and nothing above transfers automatically to a molecule that binds elsewhere. The second is jurisdictional: these sources record a United States regulatory warning and say nothing about scheduling in India, and a scheduling status recalled from memory is exactly the kind of confident specific that turns out to be wrong.

IN INDIA

The agent Indian practice must reason about, and that the standard textbook chapters do not cover, is etizolam, a thienotriazolodiazepine that behaves as a full benzodiazepine receptor agonist. The prescribing decision that follows is to treat it as a full member of this class for every purpose set out above: the same modulatory mechanism, the same subunit-determined split between sedation and anxiolysis, tolerance developing first to sedation, and a withdrawal reaction after long regular use. What must not be done is to borrow a half-life from another member to fill the gap, because the sources used here print none for it, and an assumed half-life silently decides how a reduction is planned. Its regulatory position needs the same discipline: read the current classification and the record-keeping it requires off the Drugs and Cosmetics Rules themselves rather than off a textbook.

What remains is compact and durable: one modulatory site with a ceiling built into the mechanism, a subunit map assigning sedation, anxiolysis and muscle relaxation to different parts of the same complex, a tolerance sequence beginning with the effect the patient notices, a withdrawal reaction that is the ordinary consequence of long exposure, and a dose-occupancy curve explaining why reductions get harder rather than easier.

A benzodiazepine cannot open the chloride channel without GABA, and the rest of the class follows from that: a ceiling that gives it perhaps the highest therapeutic index of all psychotropic drugs, a subunit map that splits sedation from anxiolysis, tolerance that takes sedation first, and a withdrawal reaction most long-term users will experience.

24 · Z-drugs and the other hypnotics compared

A hypnotic is compared with another hypnotic on two pharmacological axes and very little else. The first is which population of receptors the molecule prefers; the second is how long it persists in the plasma. Almost every question asked about this group is derived from those two axes: why one agent leaves impairment behind in the morning and another does not, why a drug with a weak anticonvulsant action is nonetheless an effective hypnotic, and why the withdrawal picture is so familiar. The **benzodiazepine receptor agonists** are the group in which the two axes separate most cleanly, which is why they repay being learned as pharmacology rather than as a list of names.

Two warnings about the label, because both are examined. The prefix in "non-benzodiazepine hypnotic" is a statement about chemical structure, not about mechanism and not about safety. And the letter Z is an orthographic accident: it groups agents that share a first letter and a binding site while differing in chemistry, in subunit preference and in the length of time they stay in the body. The GABA-A receptor itself, its subunit architecture and the biochemistry of inhibitory transmission belong to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes and are assumed here. Benzodiazepine class pharmacology has its own section in these notes; the use of any of these agents as a clinical treatment for insomnia

belongs to the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes, and sedative-hypnotic dependence as a disorder to the Substance use disorders notes.

about 1 hour	mean 2.4 hours	5 to 6 hours	about 6 hours
ZALEPLON, ELIMINATION HALF-LIFE	ZOLPIDEM, IMMEDIATE RELEASE	ZOPICLONE, ADULT FIGURE	ESZOPICLONE

24.1 Three chemistries, one binding site

The group is named for a letter and not for a structure. Zolpidem is an **imidazopyridine**, zaleplon a **pyrazolopyrimidine** and eszopiclone a **cyclopyrrolone**: unrelated chemical families that converge on the same pharmacological place. What unites them, on the account given in Kaplan and Sadock's Comprehensive Textbook of Psychiatry, is that they act more selectively than the classical benzodiazepines on GABA-A receptors containing the alpha-1 subunit, receptors often referred to as GABA-A1 receptors. The comparison is the load-bearing half of that sentence, and the half candidates drop. Diazepam and the other classical benzodiazepines bind with equal affinity to benzodiazepine-sensitive receptors of differing subunit composition; the Z-drugs prefer one population over the others.

That one comparative statement generates most of what follows. If a classical benzodiazepine engages every benzodiazepine-sensitive receptor more or less equally, its clinical profile will be the sum of everything those receptors do, and the adaptive consequences of occupying all of them for months will be equally broad. A molecule that prefers a single population should, on the same reasoning, deliver a narrower set of effects and adapt a narrower set of receptors. The Z-drugs are the test of that prediction, and the instructive finding is that it is borne out only in part.

RULE **The prefix records chemistry; the binding site records pharmacology.** Every drug in this group is an agonist at the benzodiazepine recognition site of the GABA-A receptor. Everything that follows from occupying that site therefore remains available: sedation, psychomotor impairment, tolerance and withdrawal on stopping. What differs between the agents is which receptor populations they prefer, and for how long they stay. Nothing in the word "non-benzodiazepine" suspends the first list.

24.2 Selectivity is graded, and it belongs to the molecule

Zolpidem is the cleanest instance. Companion to Psychiatric Studies describes it as binding selectively and with high affinity at alpha-1 containing GABA-A receptors, with roughly ten-fold lower affinity for alpha-2 and alpha-3, and none at all for alpha-5. Zaleplon sits beside it: a **full agonist** at the benzodiazepine alpha-1 receptor, which it binds with high affinity, its affinity for the alpha-2 and alpha-3 subtypes being lower. Different molecules, unrelated chemical families, one preference.

The clinical translation of that binding pattern is stated by the same source and is the part to carry away: the anxiolytic, muscle relaxant and anticonvulsant properties of zolpidem are weak. KDT Essentials of Medical Pharmacology, the Indian standard pharmacology text, makes the same attribution from the opposite direction, ascribing the relatively selective hypnotic effect of zolpidem to its preferential action at the alpha-1 subunit containing benzodiazepine receptor,

with no evident anticonvulsant, muscle relaxant or antianxiety effect to accompany it and an elimination half-life of about two hours. Independent texts, opposite framings, the same pharmacological claim.

Put those statements side by side and the mechanism becomes visible without anyone having to assert a subunit-by-subunit map of function that neither source spells out here. The actions a selective agent loses are carried by receptor populations it barely engages; the action it keeps is carried by the population it prefers. That is the entire promise of subunit selectivity, and it is why these molecules were developed.

Two riders keep the account honest, and both are examinable. Selectivity is not exclusivity: despite its preference, zolpidem is still displaced from its binding sites by **flumazenil**, which is the pharmacological proof that it is sitting where a benzodiazepine sits. And the narrowing of effect is not a narrowing of every risk. KDT records that even large doses of zolpidem do not markedly depress respiration, which is a genuine and useful contrast with the older sedatives, but respiratory margin is one property among many and it says nothing about the impairment, tolerance and withdrawal discussed below.

24.3 The cyclopyrrolones, where the tidy account breaks

The neat story survives exactly as far as the cyclopyrrolones. Companion to Psychiatric Studies states that while zopiclone is chemically a cyclopyrrolone, pharmacologically it resembles the benzodiazepines in binding indiscriminately and with high affinity across the benzodiazepine site, at alpha-1, alpha-2, alpha-3 and alpha-5 containing receptors. Kaplan and Sadock's account, quoted earlier, groups eszopiclone with zolpidem and zaleplon among the agents acting more selectively at the alpha-1 population. Both statements are in print in standard texts. They cannot be reconciled by rounding, and neither is a misprint.

The defensible reading is that selectivity here is a graded property rather than a binary one, and that it is strongest and least disputed for zolpidem and zaleplon. The cyclopyrrolones are the contested case, and a candidate is safest holding them that way: as the agents whose selectivity is claimed by one authority and denied by another, rather than as agents whose selectivity is settled. It is also worth noticing that the two statements are made about two different molecules, zopiclone in one text and eszopiclone in the other. Neither source, as quoted here, states how those two molecules are related to each other, so the reader should not assume that a finding about one transfers automatically to the other.

TRAP **Treating preference for one subunit population as the definition of a Z-drug.** Candidates learn the group as "the selective ones" and then meet a question about the cyclopyrrolones, where one standard text groups the class with the selective agents and another describes indiscriminate binding across the whole benzodiazepine site. Whichever way the paper is written, the memorised equation answers it wrongly. Learn the axis, not the slogan: selectivity is graded, it is strongest for zolpidem and zaleplon, and the cyclopyrrolones are where the sources disagree.

24.4 The second axis: how long the drug is there

Half-life does more comparative work in this group than receptor preference does. The **elimination half-life** of these agents spans a wide range for drugs given at the same point in the evening for the same purpose, and that range is what decides how much drug is still present when the patient has to function. The Maudsley Deprescribing Guidelines set the figures out drug by drug. Zaleplon is the extreme case, with an elimination half-life of approximately one hour and peak plasma concentration reached approximately one hour after oral ingestion, so that plasma levels fall to zero well before the next night-time dose is due. Zolpidem has a short but appreciably longer profile: a **mean of 2.4 hours** with a range of 0.7 to 3.5 hours, a duration of action of up to 6 hours, and peak plasma reached between 0.5 and 3 hours. The controlled release formulation produces a longer period of maximal dose and a half-life of 2.8 hours.

The cyclopyrrolones sit at the long end. Eszopiclone has an elimination half-life of **approximately 6 hours**, with peak plasma at approximately one hour. Zopiclone is the one figure the sources dispute. The Maudsley Deprescribing Guidelines give **approximately 5 hours**, with peak plasma at approximately 1.5 to 2 hours; Companion to Psychiatric Studies gives around six hours. The sensible teaching figure for an adult is therefore five to six hours, quoted with both sources named rather than resolved by picking one. The same text adds that the half-life is considerably longer in older people, at about nine hours, and that figure should be carried separately as an age-qualified statement rather than folded into the adult range. Prescribing in the elderly is treated in the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

AGENT	CHEMICAL CLASS	BINDING AT THE BENZODIAZEPINE SITE	PEAK PLASMA	ELIMINATION HALF-LIFE
Zaleplon	Pyrazolopyrimidine	Full agonist at alpha-1, high affinity; lower at alpha-2 and alpha-3	About one hour	About one hour
Zolpidem	Imidazopyridine	High affinity at alpha-1, roughly ten-fold lower at alpha-2 and alpha-3, none at alpha-5	Between 0.5 and 3 hours	Mean 2.4 hours, range 0.7 to 3.5 hours
Zopiclone	Cyclopyrrolone	Indiscriminate across alpha-1, alpha-2, alpha-3 and alpha-5, on the Companion account	About 1.5 to 2 hours	About five to six hours, sources differing
Eszopiclone	Cyclopyrrolone	Grouped with the alpha-1 preferring agents by Kaplan and Sadock	About one hour	About 6 hours

PITFALL

Reading a mean half-life as though it were the patient's own. The zolpidem figure is a mean drawn from a wide range, and the drug's stated duration of action is longer than that mean would suggest. A patient at the slow end therefore carries more drug into the morning than the textbook number predicts, and the grogginess gets attributed to poor sleep, to the illness or to anything except the tablet, on the reasoning that the half-life is short. Ask about the morning after, not only about the night.

24.5 What preference for one subunit population does not buy

If subunit selectivity delivered everything hoped of it, these agents would produce hypnosis without the residual burden of a benzodiazepine. They do not, and the evidence is at its clearest for zopiclone. Companion to Psychiatric Studies records that, like the benzodiazepines, it impairs performance motor skills; that tolerance to that impairment is only partial, so impairment continues to be evident in driving assessments with prolonged use; and that **dysgeusia**, a bitter metallic alteration in taste, is a characteristic adverse effect. The same source states that the abuse potential of the drug is clear.

The partial tolerance finding is the one that changes practice. Tolerance to sedation developing does not mean tolerance to impairment developing, and a patient who no longer feels drugged is not thereby a patient who performs normally. Abuse liability, dependence and their assessment as a disorder belong to the Substance use disorders notes; what belongs here is the pharmacological point that neither the impairment nor the abuse liability was abolished by moving away from benzodiazepine chemistry.

GOOD TO REMEMBER

Ask about driving at every review of a patient taking one of these agents, not only when the drug is started. The evidence that impairment persists comes from driving assessments in people on prolonged treatment, which is precisely the group a clinician is least likely to re-examine, because they are stable, they are not complaining, and the prescription has become a repeat. Habituation to the feeling of a drug is not evidence of recovered performance.

24.6 Receptor adaptation, and why stopping is still difficult

Here the selectivity argument does earn something measurable. Kaplan and Sadock's Comprehensive Textbook of Psychiatry reports that chronic exposure to the selective agents zolpidem and zaleplon produces less extensive change in the expression of GABA-A receptor subunits than prolonged treatment with classical benzodiazepines. That is a mechanistic result, not a marketing claim, and it is what the selectivity hypothesis predicts: fewer receptor populations occupied for months means fewer receptor populations remodelled. The important qualification travels with it. On short-term withdrawal, the changes that do occur come to resemble those seen after diazepam. Less adaptation during exposure, in other words, does not translate into a different kind of adaptation once the drug is removed.

The clinical arithmetic follows. The Maudsley Deprescribing Guidelines state that stopping these drugs, even when they have been taken exactly as prescribed, can lead to **withdrawal effects** in as many as **30 to 45 per cent** of people, and that the longer the drug has been taken and the higher the dose, the greater the risk. Two features of that statement deserve emphasis. It applies to prescribed use, so it is not a statement about misuse. And the same figure is given for each of the agents in the group, so it is not a peculiarity of one molecule that a switch within the class would solve.

TRAP **Reading "non-benzodiazepine" as "not dependence forming", and stopping abruptly on that basis.** The withdrawal frequency quoted above is not a benzodiazepine figure borrowed for effect; it is the figure given for these drugs, in people who took them as prescribed. The examination version of this error is the candidate who lists absence of dependence liability as an advantage of the class. The clinical version is the prescriber who discontinues a long-standing hypnotic without a plan, because the label said the drug was not a benzodiazepine.

24.7 The hypnotics that never touch the benzodiazepine site

The benzodiazepine receptor agonists are one family among several, and a comparison confined to them misses the axis that separates the classes. Practical Psychopharmacology sets out the wider field: beyond the benzodiazepine receptor agonists, the hypnotic families in use are the antihistamines, the **melatonin receptor agonists** and the **orexin receptor antagonists**. Read as a list of targets rather than a list of drugs, those are distinct ways of producing sleep.

- **GABA-A modulation.** The benzodiazepines and the Z-drugs, acting at the benzodiazepine recognition site, with the profile set out above.
- **Histamine blockade.** The antihistamines, several of which are available without prescription in many countries, producing sedation as an extension of a receptor action that is otherwise an unwanted effect.
- **Melatonin agonism.** Agents acting at melatonin receptors, working with the timing system rather than by depressing arousal.
- **Orexin antagonism.** Agents blocking orexin signalling, and therefore removing a wake-promoting drive rather than adding an inhibitory one.

That last distinction is the conceptual pay-off. A GABA-A modulator produces sleep by increasing inhibition; an orexin antagonist produces it by subtracting wakefulness; a melatonin agonist works through the circadian signal. Because the two newer families act nowhere near the benzodiazepine recognition site, the tolerance and withdrawal pharmacology set out above does not transfer to them by the same mechanism. What the sources quoted here do not supply are half-life or receptor-occupancy figures for the melatonin and orexin agents, so no kinetic comparison with the Z-drugs can honestly be printed, and none is attempted. The licensing position of individual agents also differs between countries, so the class names travel where product names do not. Their place in the treatment of insomnia belongs to the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes.

24.8 Comparing any two hypnotics

The comparison that examinations reward, and the one that survives the arrival of new molecules, runs in a fixed order. Establish the mechanism family first, because it decides whether the dependence and withdrawal pharmacology of the benzodiazepine site is in play at all. Within that family, establish the subunit preference, remembering that it is graded and that for the cyclopyrrolones it is disputed. Then take the half-life, which predicts how much drug is present when the patient must drive, work or supervise a child. Only then read the adverse-effect signature, which for this group includes psychomotor impairment that tolerance does not fully abolish and, for one agent, a characteristic taste disturbance.

MNEMONIC

Zaleplon, Zolpidem, Zopiclone, eszopiclone: shortest half-life first

The agents listed in ascending order of elimination half-life are also the order in which they are usually asked about, and that order predicts, other things being equal, how much drug is still present when the alarm goes off. **Za**leplon is the outlier at roughly an hour, **Zol**pidem is short but longer, and the cyclopyrrolones sit at the long end.

What none of this settles is whether a given patient should be taking a hypnotic at all. That decision belongs with the assessment of the sleep complaint. The pharmacology is what makes it executable: it tells the prescriber what the chosen molecule will still be doing at breakfast, what will happen when the patient stops it, and which of the reassuring things said about the class are properties of the drug rather than properties of its name.

The Z-drugs are benzodiazepine-site agonists whose prefix hides the fact: preference for the alpha-1 population is graded and strongest for zolpidem and zaleplon, half-life decides what is left in the morning, and withdrawal on stopping is common in people who took the drug exactly as prescribed.

25 · The non-benzodiazepine anxiolytics and the cognitive enhancers

Every drug gathered here is named for what it is not. The non-benzodiazepine anxiolytics are defined by a negative: they relieve anxiety without acting where a benzodiazepine acts, each by a different molecular route, so the group shares no receptor, no time course and no adverse-effect signature. The cognitive enhancers are the same kind of grouping seen from the other side, three drugs that inhibit an enzyme and one that plugs an ion channel, sharing a shelf rather than a mechanism. Treat either family as a single class and most questions about it become unanswerable; read each agent by its target and the answers fall out of the target.

The marks sit on dissociations, and three are worth naming in advance: a plasma half-life of a few hours against an onset measured in weeks; a plasma half-life that says nothing about how long an enzyme stays inhibited; and one mechanism the standard textbooks label with two different words. Whether any of these agents is the right treatment for a particular anxiety disorder belongs to the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes, and their

indication, staging and efficacy in dementia belong to the Dementias and neurocognitive disorders notes. Cross-class interaction tables and monitoring schedules belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

2 to 4 hours	10 days	1.6 hours	about 70 hours
BUSPIRONE PLASMA HALF-LIFE	BEFORE ITS ANXIOLYTIC EFFECT BECOMES APPARENT	RIVASTIGMINE MEAN ELIMINATION HALF-LIFE	DONEPEZIL HALF-LIFE, HENCE ONCE-DAILY DOSING

25.1 An azapirone, and the receptor it never touches

Buspirone is an **azapirone**, also written as an azaspirodecanedione, and the single most important fact about it is a negative one: it does not act on the GABA-A benzodiazepine chloride ionophore at all. That one statement generates most of what follows, because every property a benzodiazepine derives from modulating a chloride channel is one buspirone has no route to.

Where it does act is on a serotonin receptor. Buspirone is a **partial 5-HT1A agonist** and a relatively weak dopamine antagonist, apparently at predominantly presynaptic sites. The dopaminergic component is worth holding separately: it is the weaker action, and its site is qualified rather than asserted. What partial agonism does to a prescribing decision is developed where receptor theory is treated in this chapter, and the underlying receptor biology sits in the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

Two cautions belong with that receptor address. The Companion to Psychiatric Studies, having given the profile, describes buspirone's complex actions as incompletely understood, so the profile is a location and not a completed explanation. And candidates who have learnt that anxiolysis means a chloride channel will translate a serotonergic anxiolytic back into GABAergic language, then predict sedation, anticonvulsant activity and cross-tolerance that this drug does not have. The mechanism is not a paraphrase of the benzodiazepine mechanism; it is a different one that happens to earn the same therapeutic label.

25.2 A short half-life against a slow onset

The kinetics are unremarkable, and worth learning because they explain so little. Buspirone is rapidly absorbed, undergoes extensive presystemic metabolism and has a half-life of only **two to four hours**. A drug cleared that quickly rises and falls several times a day, which on the reasoning used everywhere else should predict a rapid, easily titrated effect.

It predicts nothing of the sort. Buspirone's action is not instantaneous, requiring some 10 days to become apparent and up to 3 weeks to develop fully. Set those two figures side by side and the arithmetic is unavoidable: the effect cannot be a simple function of plasma concentration, because the concentration has completed hundreds of cycles before the effect appears. Whatever produces anxiolysis here is downstream of receptor occupancy and is timed like an antidepressant response, not like a sedative.

The exact figure is not agreed. Tasman's psychiatry text puts the effect at least two weeks away and reports clinical experience suggesting three to four weeks before a truly beneficial effect is seen, while the Companion gives the shorter figure above. Nothing turns on which is right: the

teachable envelope runs from about ten days to about four weeks, and the contrast that earns the mark, against the very rapid onset of a benzodiazepine, survives either.

One vulnerability follows from the presystemic step. Buspirone's metabolism is inhibited by naringenin in grapefruit juice, which produces substantial increases in blood levels. That is a class-level property rather than an interaction table; the tables belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

TRAP **Quoting a single onset figure for buspirone.** Two standard texts give different numbers, and an answer key built on either will mark the other wrong. Say that the effect takes days to weeks rather than minutes, give the envelope, and name the contrast with the benzodiazepine; that answer is correct against both sources and is what the question is actually testing.

25.3 One effect where a benzodiazepine has four

Buspirone's spectrum of action is restricted to anxiety, with little or no hypnotic, anticonvulsant or muscle relaxant component. A benzodiazepine carries all four actions on one molecule because all four descend from the same channel; buspirone carries one, because it is not at that channel. The practical corollary is the hardest single rule here: it should not be used to substitute for benzodiazepines or alcohol in withdrawal regimes, and the same text warns that it has a mild proconvulsant action and should be avoided in those at risk of fits and in combination with monoamine oxidase inhibitors. Benzodiazepine dependence as a disorder belongs to the Substance use disorders notes.

That same absence of agonist activity puts buspirone at the far end of the misuse-liability spectrum. Buspirone is essentially devoid of benzodiazepine agonist activity and has little or no misuse liability in the predictive human models used to rank sedative-hypnotic drugs, placing it below the benzodiazepines and far below methaqualone and the barbiturates. Read the qualification with the claim: those rankings come from laboratory paradigms of drug liking and from the ARCI-MBG scale, not from field data, so what has been shown is an absence of reinforcing subjective effect, not an absence of misuse in the world.

Stopping the drug behaves accordingly. Discontinuation of buspirone has not been associated with the development of withdrawal symptoms, although patients discontinued from pharmacotherapy tend to relapse over time. That pairing is the honest version of the class advantage: relapse and withdrawal are different events, and a drug can be free of one while remaining fully subject to the other.

PITFALL

Starting buspirone to cover a benzodiazepine taper, then reading the patient's distress as anxiety returning. A drug with no action at the benzodiazepine site cannot occupy the receptor the patient is withdrawing from, so the taper runs unopposed while an agent that will not work for another two weeks is credited with holding it. The error is attractive because the two drugs share the word anxiolytic, and the proconvulsant caution makes it worse, since the population being tapered is the population with a lowered seizure threshold.

25.4 Beta-adrenergic blockade, the gabapentinoids, and onset as the axis

The other non-benzodiazepine routes are narrower than their reputation. Beta-adrenergic blockers such as propranolol may achieve short-term control of tremor and palpitations, which can be the most handicapping symptoms of specific social phobias such as performance anxiety among musicians, and it is doubtful whether they are more generally effective in social phobia. That sentence is the whole of what the Shorter Oxford Textbook of Psychiatry supports: two target symptoms, one narrow situation, and no mechanism at all.

Two things commonly taught alongside it are absent from the sources here, and both should be omitted rather than reconstructed. The familiar division of anxiety into somatic symptoms that a beta-blocker treats and psychic symptoms that it does not is not in this passage. Neither is the lipophilicity argument, the claim that propranolol is preferred over atenolol because it enters the central nervous system: no source here supports choosing a beta-blocker for its degree of central penetration, so that reasoning should not be offered as the pharmacological basis for the choice.

The gabapentinoids are the third route and the most often misfiled. Pregabalin binds to a subunit of voltage-gated calcium channels in the central nervous system, is unusually not protein bound and undergoes virtually no metabolism, so its interactions are few, and the Companion is explicit that how that binding relates to the drug's mode of action is unknown. The subunit is named elsewhere: **the alpha-2-delta protein** of the voltage-sensitive calcium channel is the target of pregabalin and gabapentin, and that protein may regulate conformational changes that alter how the channel opens and closes. The name gabapentinoid records a chemical resemblance, not a receptor: there is no action at the GABA-A receptor here, and a candidate who assumes one predicts the wrong interactions and the wrong withdrawal.

On a single axis of onset the group organises itself, and the axis is mechanism, not indication.

AGENT OR CLASS	MOLECULAR TARGET AS STATED	ONSET OF ANXIOLYTIC EFFECT
Buspirone	5-HT1A partial agonist	Delayed, at about two weeks
Beta-adrenergic blockers	Blockade of beta-adrenergic receptors	Rapid
Pregabalin	GABAergic by a mechanism not fully known	Rapid
Benzodiazepines	Modulation of the GABA-A receptor	Rapid to very rapid

25.5 The cognitive enhancers: two mechanisms, not one class

The second family is smaller than it looks and splits cleanly in two. Three cholinesterase inhibitors are currently available in the United States: donepezil, rivastigmine and galantamine, with tacrine, the first agent approved, no longer on that market. Standing apart from them, memantine is an NMDA receptor antagonist approved for moderate to severe Alzheimer disease. Both approval statements are specific to the United States and should be quoted as such. Which agent is indicated at which stage, and what either family achieves clinically, belongs to the Dementias and neurocognitive disorders notes.

The cholinergic three share a target that is more divisible than the class name suggests. Two enzymes hydrolyse acetylcholine, **acetylcholinesterase** and **butyrylcholinesterase**, and the cholinesterase inhibitors differ in the extent to which they selectively inhibit each. The second enzyme is the minor partner centrally, making up about 10 per cent of the total cholinesterase in the central nervous system according to the same source, which is why a drug that also blocks it is doing something additional rather than something more.

That split is what makes the within-class comparison tractable: there is no single quantity called cholinesterase inhibition to rank these drugs on, only two enzymes, a duration of binding, a nicotinic action that one agent has, and a set of kinetics.

25.6 What actually separates the three cholinesterase inhibitors

The enzyme selectivities are the first column and they are stated cleanly. Donepezil inhibits acetylcholinesterase with higher selectivity, rivastigmine inhibits both acetylcholinesterase and butyrylcholinesterase, and galantamine selectively inhibits acetylcholinesterase while also producing allosteric modulation of nicotinic receptors. The withdrawn agent completes the range in the opposite direction: tacrine inhibited butyrylcholinesterase 18 times more potently than acetylcholinesterase, which shows that selectivity within this class is a real variable.

Galantamine's second action is a genuinely different mechanism and is examined as such. **Allosteric modulation of nicotinic receptors** is distinct from the action of acetylcholine and from that of nicotinic agonists or antagonists, and it upregulates or potentiates nicotinic ion channel activity in response to acetylcholine, so cholinergic transmission is enhanced by a route that never passes through the enzyme.

The second axis is time on the enzyme rather than affinity for it. Cholinesterase inhibitors are classified as reversible, pseudo-irreversible or irreversible according to how long they bind the enzyme, and the three arms are worth separating explicitly.

- Reversible inhibitors bind the enzyme for seconds to minutes, so inhibition tracks the plasma concentration closely.
- **Pseudo-irreversible inhibitors** bind for hours, so inhibition outlasts the concentration that produced it.
- Irreversible inhibitors sit at the far end of the same axis, where the enzyme is not available again until new enzyme has been made.

MNEMONIC

Donepezil one, rivastigmine both, galantamine plus nicotinic

Three agents, three answers to one question about the target. **Donepezil** acts on acetylcholinesterase preferentially. **Rivastigmine** takes both enzymes. **Galantamine** is selective for the same enzyme as donepezil but adds a nicotinic action nobody else in the class has. The line records what each one does, never which is better.

25.7 Kinetics, and why the plasma half-life is the wrong question

The elimination half-lives inside this small class span a very wide range, and each must be read against the binding classification rather than on its own. Rivastigmine is rapidly absorbed orally with a mean time to maximum plasma concentration of one hour, range 0.8 to 1.7 hours, and has a mean elimination half-life of **1.6 hours**. Read naively, that predicts a drug present for only part of the day. It does not, because rivastigmine's inhibition is pseudo-irreversible, and for such an inhibitor the plasma half-life and the duration of enzyme inhibition are different quantities that need not resemble each other.

Galantamine sits in the middle and carries the class's one metabolic complication. Galantamine immediate release has an elimination half-life of about seven hours, and galantamine is mainly metabolised by hepatic CYP2D6 and CYP3A4, with the CYP2D6 metabolite sanguinine accounting for 20 per cent of the parent drug and being three times more potent than the parent drug at inhibiting acetylcholinesterase. An active metabolite more potent than its parent, formed by a polymorphic enzyme, is the classic setting for a genotype-driven difference in effect, though the source's own position is that the poor-metaboliser difference is probably not significant given the large variation between patients. The enzymology sits in the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes, and the cross-class interaction consequences in the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

The two long-acting agents anchor the other end. Donepezil's schedule is derived from its kinetics: the Indian pharmacology standard attributes its once-daily administration at bed time to a long half-life of about 70 hours, a distinct advantage over rivastigmine and galantamine, both of which need twice-daily dosing, and the Comprehensive Textbook of Psychiatry gives the same order of half-life for the same reason. Memantine is longer again: it reaches peak plasma concentration within three to seven hours, has an elimination half-life of **between 60 and 80 hours**, and is largely excreted unchanged in the urine, the remainder being converted to three minimally active polar metabolites, with steady state reached separately at about 21 days.

AGENT	ENZYME OR RECEPTOR ACTION	ELIMINATION HALF-LIFE	KINETIC CONSEQUENCE
Donepezil	Inhibits acetylcholinesterase with higher selectivity	About 70 hours	Once daily at bed time
Rivastigmine	Inhibits both acetylcholinesterase and butyrylcholinesterase	1.6 hours, mean	Twice-daily dosing
Galantamine	Selective for acetylcholinesterase, plus nicotinic allosteric modulation	About seven hours, immediate release	Twice-daily dosing
Memantine	NMDA receptor antagonist	60 to 80 hours	Steady state at about 21 days
Tacrine, withdrawn	Inhibited butyrylcholinesterase 18 times more potently than acetylcholinesterase	Not stated in these sources	Withdrawn from the United States market

GOOD TO REMEMBER

A half-life answers how often the drug must be given and how long it takes to reach steady state. It does not answer how long the target stays engaged, which for a pseudo-irreversible enzyme inhibitor is a separate measurement. With an agent whose steady state is three weeks away, anything assessed in the first fortnight measures the titration rather than the medicine.

RULE **A short half-life buys neither a fast onset nor a short action.** Buspirone is cleared in a few hours and takes weeks to work; rivastigmine is cleared faster still and inhibits its enzyme for longer than it is present. Onset is set by whatever process actually mediates the effect, and offset by how long the target stays engaged. The plasma curve sets the dosing interval and nothing else.

25.8 Memantine, and the block that works only while the channel is open

Memantine is the one agent here that is not a cholinergic drug, and its mechanism is the most precisely specified claim in the section. It is an **uncompetitive open-channel NMDA receptor antagonist**, and Stahl's account gives the four properties together: low to moderate affinity, voltage dependence, and fast blocking and unblocking kinetics, with the drug working at the same site as magnesium, plugging the ion channel where the magnesium ion normally blocks it at rest. Each clause does work. Because the drug occupies the channel rather than the glutamate recognition site, it can bind only once the channel has been opened, which is what voltage dependence means here; and because affinity is low and the kinetics fast, a barrage of glutamate from normal neurotransmission readily and quickly reverses the block.

The therapeutic logic follows directly. Chronic blockade of the NMDA receptor would hypothetically interfere with memory formation and neuroplasticity, so a drug meant to damp excessive tonic activation without abolishing phasic signalling has to be this kind of antagonist

and no other. Tasman's text describes the same mechanism independently, as an uncompetitive antagonist binding the NMDA receptor-operated cation channels and acting by inhibiting calcium influx.

One caution about the literature is worth carrying into the examination hall. At least one printing of a major psychiatry textbook carries a sentence that reads as though memantine blocked glutamate binding at the receptor. That is incompatible with open-channel block and with the surrounding paragraph, which describes selective uncompetitive blockade and strong voltage dependence. Where one sentence contradicts the mechanism its own paragraph describes, take the paragraph.

TRAP **Treating uncompetitive and noncompetitive as a factual disagreement about memantine.** Every source that spells the mechanism out in prose calls the drug uncompetitive, and the same books print noncompetitive in figure legends and passing mentions. The mechanism described is identical in both: binding inside the channel at the magnesium site, only while the channel is open, with low affinity and fast kinetics. Learn the mechanism, use uncompetitive, and treat a mismatch against an answer key on this one word as a labelling difference rather than an error of fact.

25.9 Nootropic: a claim about a category, not a mechanism

The word that ends this section is not a pharmacological description at all. Piracetam is a cyclic GABA derivative with no GABA-like activity, and the word **nootropic** denotes a drug said to selectively improve the efficiency of higher telencephalic integrative activities. Both halves repay attention. The molecule resembles a transmitter it does not act like, which is the trap the gabapentinoid name sets earlier; and the label names a hoped-for functional outcome rather than a target, so nothing about a receptor, an enzyme or a channel can be derived from it. A category defined by its intended effect is a claim awaiting evidence, not a mechanism.

The evidence has been examined and the verdict is unambiguous. Piracetam has long been promoted in India and some other countries for cognitive impairment and dementia in the elderly and for intellectual disability in children, but a Cochrane Database review of 2004 concluded that published data do not support such use, and in the United Kingdom it is approved only as adjunctive treatment of cortical myoclonus. The narrow surviving approval is the tell: the demonstrable effect was licensed, the broad cognitive claim that named the category was not.

IN INDIA

This is the one agent here where the prescribing decision is made against local practice rather than with it. Piracetam is promoted in this country for exactly the two populations in which the evidence was examined and found wanting, so it should not be started as a cognitive enhancer for an older adult with cognitive impairment or for a child with intellectual disability, and a request to add it to an existing regimen has no pharmacological argument behind it.

When a patient arrives already taking it, ask which claim it was started on. If the answer is cognition, say plainly that the licensed indication surviving review is a movement disorder rather than a cognitive one. The drug is well known in Indian households, and quietly stopping the tablet reads as a treatment being taken away.

Nothing in this section shares a mechanism with the drug it is offered as an alternative to, so nothing in it inherits that drug's speed, its spectrum or its behaviour on withdrawal.

26 · Psychostimulants and atomoxetine: onset, duration and diversion

What separates one stimulant preparation from another is almost entirely kinetic. The molecules in routine use are few: methylphenidate, dexamfetamine and the mixed amphetamine salts all raise synaptic catecholamines at the monoamine transporters. What differs between the products, and what is examined here, is rate. How fast the drug arrives, how long its effect lasts and how quickly it occupies a transporter determine both the shape of the therapeutic day and the probability that the tablet is swallowed by somebody other than the patient. Atomoxetine belongs in the same section for the opposite reason: it is not a stimulant at all, and set against the stimulants it demonstrates that abuse liability is a property of where and how fast a transporter is occupied rather than of which transmitter is raised. The mechanism itself, reuptake blockade against transporter-mediated release, and the indication-specific dosing, titration and adverse-effect monitoring of these agents, belong to the ADHD and specific learning/communication disorders notes; stimulant misuse as a disorder belongs to the Substance use disorders notes.

within 30 minutes	4 to 6 weeks	up to 12 hours	15 per cent
ORAL STIMULANT ONSET OF ACTION	ATOMOXETINE TIME TO EFFECT	LISDEXAMFETAMINE DURATION OF ACTION	METHYLPHENIDATE PLASMA PROTEIN BINDING

26.1 The one property the whole class is organised around

The class begins at the gut wall. Tasman Psychiatry puts it simply: psychostimulants are rapidly absorbed from the gut and therefore act quickly, an oral dose beginning to work **within the first 30 minutes** of ingestion. Almost nothing else in psychiatric pharmacology behaves like this: an antidepressant or an antipsychotic is judged over weeks, a stimulant before the patient has left the waiting room.

Three consequences organise the rest of the topic. First, the therapeutic trial is short, so response and non-response are both visible quickly rather than inferred from a plasma level. Second, because the effect appears and disappears within a single day, the clinically interesting number is

not the elimination half-life but the **effect duration**, the length of time a single administration continues to do useful work. Third, a drug that reaches the brain fast enough to be felt can by that same property reinforce its own consumption. Speed is at once the therapeutic advantage of the class and the reason it is regulated.

26.2 Onset, peak, half-life and duration are four different numbers

Residents lose marks by collapsing four independent parameters into one. Tasman Psychiatry tabulates them separately for the immediate-release preparations.

IMMEDIATE-RELEASE PREPARATION	EFFECT ONSET (MINUTES)	ELIMINATION HALF-LIFE (HOURS)	EFFECT DURATION (HOURS)
Methylphenidate	30 to 60	2 to 4	3 to 5
Dexamfetamine	30 to 60	3 to 5	4 to 6
Mixed amphetamine salts	30	3 to 5	4 to 6

The time to peak concentration is the fourth parameter and it separates the preparations least: the same table gives a time to peak of one to three hours for methylphenidate, for dexamphetamine and for the mixed salts alike. What does separate them is the pair at the right of the table. Methylphenidate has both the shortest elimination half-life, at **two to four hours**, and the shortest effect duration; the amphetamine preparations sit a step longer on both axes. That within-class ranking is the durable teaching point, and it is why immediate-release methylphenidate has needed more frequent administration than an amphetamine to cover the same waking day. Notice too that duration is not read off the half-life: for methylphenidate the tabulated effect outlasts the tabulated half-life, as you would expect of a drug whose action depends on transporter occupancy rather than plasma concentration. A half-life predicts how long a drug takes to leave the body, not how long the patient feels better, and here the second question is the clinical one.

TRAP **Quoting one methylphenidate half-life as though the standard texts agreed, and then treating that figure as the duration of action.** Two errors travel together. The first is that Tasman Psychiatry is internally inconsistent about the elimination half-life: its parameter table gives two to four hours, while the prose of the same chapter reports that methylphenidate peaks in plasma at 2 to 2.5 hours and falls to half the peak after 3 hours. This is one book disagreeing with itself, and because the prose point estimate sits inside the tabulated range, the range is the safer figure to carry into an answer. Declare the discrepancy rather than resolving it silently. The second error is the more consequential: candidates who have memorised a short half-life conclude that every methylphenidate product needs several administrations a day, which is false the moment a modified-release preparation is prescribed. Half-life belongs to the molecule; duration of action belongs to the product.

26.3 Where methylphenidate actually goes

Methylphenidate has an unusual disposition for a psychotropic, and its interaction profile falls straight out of it. The parent compound is hydrolysed at its ester group to yield **ritalinic acid**, an inactive product, and approximately 20 per cent is oxidised by the liver to secondary metabolites.

The dominant route is an esterase reaction rather than a cytochrome one, which is why this drug does not sit at the centre of the psychotropic substrate map that the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes set out.

The second peculiarity is how little of it needs to be there. Its low plasma protein binding, of about 15 per cent, makes it highly available to cross the blood-brain barrier, which is part of why relatively low plasma concentrations are effective. The comparison the source draws is striking: clinically effective oral methylphenidate produces plasma concentrations as low as 7 to 10 nanograms per millilitre, against dexamfetamine concentrations between 40 and 120 nanograms per millilitre in children with ADHD. Such a gap suggests a large **first-pass effect**, and the reported oral bioavailability has itself moved: once given in the 80 per cent range, it is placed by more recent studies close to 30 per cent. A drug largely destroyed before it reaches the circulation, which nonetheless works an order of magnitude below its comparator, is one whose plasma concentration tells you very little.

The site of that metabolism is contested inside a single book, and the honest position is to say so. Tasman Psychiatry states in prose that methylphenidate is primarily metabolised extrahepatically by **de-esterification** to an inactive metabolite, while the parameter table of the same chapter records its metabolism as hepatic, and the passage above assigns a fifth of the drug to hepatic oxidation. Teach the mechanism, not the compartment: clearance is mainly by de-esterification to inactive ritalinic acid with a minor oxidative route, and the extrahepatic claim is not settled enough to offer as though it were.

26.4 The release mechanism, not the molecule, sets the duration

Formulation is the principal therapeutic variable in this class, and it is engineering rather than chemistry. The same molecule can be built to act for a third of a working day or for the whole of one, and the design that achieves it is stated on the product, not deduced from the drug. Tasman Psychiatry maps the **release mechanism** of each preparation onto the duration it buys.

PREPARATION	MECHANISM OF DRUG RELEASE	DURATION OF ACTION (HOURS)
Mixed amphetamine salts	Single pulse	4 to 6
Mixed amphetamine salts, long acting	Two pulses	8 to 10
Methylphenidate, sustained release	Wax matrix	3 to 8
OROS methylphenidate	Three pulses	10 to 12
Transdermal methylphenidate patch	Transdermal	10 to 12
Lisdexamfetamine	Delayed single pulse	Up to 12

Read down the duration column. One pulse delivers a single bolus and buys the shortest cover; two pulses buy roughly double; three pulses, or a transdermal system that bypasses the gut and releases continuously, buy the longest. The wax matrix is the outlier: its band is the widest of any

design here, so such a product is the least predictable and the least suitable when the timing of cover matters. The prodrug reaches a long duration by another route, its delay being metabolic rather than physical.

PITFALL

Switching a patient between two preparations of the same drug on total daily dose alone. Methylphenidate is not one product but several, and the tabulated durations of a wax-matrix and an osmotic three-pulse preparation do not overlap at their extremes. A switch made on the drug name changes the shape of the day even when the total dose is unchanged, and the complaint that follows is usually loss of afternoon efficacy or a new evening problem, neither a failure of the molecule. Name the preparation, not just the drug, on every prescription and in every handover note.

26.5 Why speed, and not the molecule, is the abuse variable

The pharmacology of diversion is the core of this topic, and it is a story about rate. Tasman Psychiatry records that the slow reversal of methylphenidate binding to the dopamine transporter has been interpreted to mean that methylphenidate is not as reinforcing as cocaine, and is therefore less likely to lead to self-administration, even though therapeutic oral doses significantly increase extracellular dopamine. That last clause is the part candidates miss. The dopamine rises; reinforcement does not follow, because it tracks how abruptly occupancy changes rather than how much transmitter is present.

Stahl's Essential Psychopharmacology generalises the principle into a class statement, and flags it throughout as hypothetical. Extended-release oral stimulants, the transdermal methylphenidate patch and the prodrug lisdexamfetamine are considered **slow-dose stimulants**: three quite different pieces of engineering converging on one kinetic result.

- an oral extended-release system, which meters the drug out of the tablet;
- a transdermal patch, which meters it through the skin and never presents a gut bolus at all;
- a prodrug, which meters it by requiring an enzymatic step before any active drug exists.

The shared mechanism is regional. These preparations occupy noradrenaline transporters in the prefrontal cortex with slow enough onset, and for long enough duration, to enhance tonic signalling, but they do not occupy dopamine transporters in the nucleus accumbens quickly or extensively enough to increase **phasic signalling** at D2 receptors, and it is that failure to drive the accumbens phasically that hypothetically reduces abuse potential. Tonic cortical signalling is what the therapeutic effect is thought to need, phasic accumbens signalling is what reinforcement needs, and formulation supplies the first while withholding the second.

Lisdexamfetamine is the purest expression of the idea. It is a **prodrug** in which dexamfetamine is complexed with the amino acid lysine and is inactive in that form, and it is broken down in red blood cells so that dexamfetamine is gradually made available, which is why it is unlikely to be abused for recreational or dependency-driven purposes. Because the rate-limiting step lives inside the erythrocyte rather than the capsule, crushing, snorting or injecting the contents

defeats nothing: the route changes, the conversion step does not. The source applies the same practical logic to extended-release methylphenidate preparations.

RULE **Abuse liability in this class is set by the rate and the site of transporter occupancy, not by the identity of the drug or by the amount of dopamine released.** Every fact in this section applies that one sentence. Therapeutic oral methylphenidate raises extracellular dopamine and is still poorly reinforcing, because its binding to the transporter reverses slowly. A modified-release and an immediate-release preparation of the same molecule carry different risk, differing in nothing but rate. A prodrug is safer than its own active metabolite, because a conversion step is inserted between the mouth and the accumbens.

MNEMONIC

Slow in, long on, cortex not accumbens

Slow in: the onset must be gradual, which is what every extended-release, transdermal and prodrug design is engineered to achieve. **Long on**: occupancy must be sustained rather than spiky, because tonic signalling is what the therapeutic effect requires. **Cortex not accumbens**: the transporters to occupy are the noradrenaline transporters of the prefrontal cortex, not the dopamine transporters of the nucleus accumbens.

26.6 The regulatory expression of a kinetic fact

Regulation follows the pharmacology, and the examinable content is the reasoning rather than the schedule. The Maudsley Prescribing Guidelines record that both methylphenidate and dexamfetamine are Schedule 2 Controlled Drugs in the United Kingdom, that prescriptions should be written appropriately, with the total amount in words and figures, and that they should cover **a maximum supply of 30 days**. The same source makes a within-class discrimination worth carrying: dexamfetamine is probably more likely than methylphenidate to be diverted and misused. The source states the ranking without giving a mechanism, so carry it as an empirical comparison rather than as something deduced from the kinetics above, and comparative rather than a licence to regard methylphenidate as safe.

The scheduling quoted here is one jurisdiction's and should be labelled as such whenever it is used. The corresponding Indian position, the schedule these drugs occupy in Indian law and the prescription requirements that follow, is not established in the sources behind this section and is deliberately not asserted; take it from current Indian statute and current drug-controller requirements, which change more often than pharmacology does. What transfers across jurisdictions is the underlying claim: a preparation is regulated because of how fast it reaches the brain, so a preparation engineered to reach the brain slowly is making a regulatory argument as well as a therapeutic one.

26.7 Atomoxetine: the negative control for the whole argument

Atomoxetine is best taught as the experiment that isolates the variable. KDT Essentials of Medical Pharmacology defines it precisely: a **selective noradrenaline reuptake inhibitor**, unrelated to amphetamine and to imipramine, which does not enhance dopamine release in the brain, and is neither a central nervous system stimulant nor an antidepressant. Every element of

that definition works. It is not an amphetamine derivative, so structural analogy predicts nothing. It is not a tricyclic, so that class's receptor promiscuity does not apply. And it is not an antidepressant despite inhibiting a monoamine transporter, a useful reminder that mechanism does not confer class membership.

How it avoids being reinforcing is a matter of regional anatomy. Stahl's Essential Psychopharmacology explains that by blocking noradrenaline transporters, atomoxetine raises both noradrenaline and dopamine in the prefrontal cortex, where the inactivation of both transmitters is largely by those transporters, while the relative lack of noradrenaline transporters in the nucleus accumbens prevents it from raising either transmitter there, and that regional difference reduces the risk of abuse. This is the slow-dose argument by another road: the stimulant preparations avoid the accumbens by being slow, atomoxetine because the transporter it blocks is scarce there. Other noradrenaline transporter inhibitors would be expected to have the same effects.

Its disposition introduces the topic's one pharmacogenetic variable. Atomoxetine is absorbed orally, hydroxylated by CYP2D6 and excreted in the urine mainly as the glucuronide, and while the majority of individuals are extensive metabolisers, a few are **poor metabolisers** because of polymorphism of CYP2D6. The enzymology belongs to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes, and co-prescription with a CYP2D6 inhibitor to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes. What belongs here is the class-level consequence: a drug cleared through a polymorphic enzyme has a wider spread of exposure across a population than one cleared by an esterase.

The clinical headline is the timescale. The Maudsley Prescribing Guidelines give the time to effect for atomoxetine as **approximately four to six weeks**, and note in the same breath that it is not a controlled drug and may be useful where the diversion of stimulants is a problem. Held against an oral stimulant working inside the first half hour, that is an enormous gap in time to effect between two drugs given for the same purpose, and it is the comparison most likely to be examined here.

TRAP **Treating atomoxetine as a slow stimulant, and abandoning it early because nothing has happened.** The stem describes a patient started on atomoxetine who is unchanged after a fortnight, and offers a switch or a dose increase as the tempting answer. Both are premature. A stimulant declares itself the same day, so an unhelpful stimulant trial is over quickly; a noradrenaline reuptake inhibitor is assessed over weeks, and the quoted interval has not elapsed. The deeper version is classificatory: candidates who file atomoxetine under stimulants expect it to carry stimulant liability, and are then unable to explain why it is the agent recommended where diversion is the problem.

GOOD TO REMEMBER

Three questions convert this section into a prescribing conversation. What does the day need to look like, which is a question about effect duration and therefore about the release mechanism rather than the drug name. What happens to the tablets between the clinic and the patient, which kinetics answers: a prodrug or a patch where supervised immediate-release dosing cannot be guaranteed. And how long is everybody prepared to wait, because a family told that timescale in advance behaves very differently from one told nothing.

In this class the kinetics are the pharmacology: onset, effect duration and half-life are separate numbers that must never be substituted for one another, the release mechanism rather than the molecule sets how long a preparation works, and both the therapeutic effect and the abuse liability follow from how fast and where transporter occupancy changes, which is why a prodrug, a patch and a non-stimulant reuptake inhibitor all reduce diversion risk by the same argument.

Consolidate and apply

27 · Worked vignettes

A pharmacology vignette asks what the drug is doing, not what the diagnosis is. The examiner rarely asks a candidate to define a receptor. Far more often a short stem describes a patient already on treatment, or about to start it, plants one pharmacological fact, and waits to see whether the candidate reasons from it. That planted fact is almost always one of four things: a binding profile, an elimination half-life, an enzyme that some other drug has just inhibited or induced, or a margin between the concentration that works and the concentration that harms. This closing section teaches no new agent and no new class. It rehearses the pharmacology assembled across this chapter as a reading skill, moving from a paragraph of history to the mechanism that explains it.

The discipline that earns marks here is resisting the reflex to answer with a treatment algorithm. A stem that describes a patient whose sedation is intolerable is not asking which antipsychotic treats schizophrenia; it is asking which receptor family the current agent occupies and what a within-class move would change. Where a stem genuinely asks which drug to choose for a disorder, the reasoning belongs to that disorder's own chapter, and the antipsychotic and antidepressant selection decisions are set out in the Schizophrenia: management, antipsychotics and adverse effects notes and the Mood disorders: pharmacotherapy notes respectively. What this chapter supplies, and what a good vignette answer uses, is the layer underneath those decisions.

27.1 Four questions read almost any pharmacology stem

Before deciding what a stem is asking, run the same four questions across it. They are the four axes on which psychotropic agents actually differ, and between them they account for most of the

discriminating information an examiner can plant in a short paragraph.

Binding	Duration	Company	Margin
WHAT IT OCCUPIES	HOW LONG IT LASTS	WHAT ELSE IS ON BOARD	HOW MUCH ERROR IT FORGIVES

Binding asks which receptors and transporters the agent touches and with what relative affinity, because a drug's unwanted effects are largely the affinities that were not the point of prescribing it. Duration asks how long the drug persists, which governs how quickly it accumulates, how forgiving it is of a missed dose, and how quickly anything at all can be judged. Company asks what else the patient is taking, including the drugs a psychiatrist did not prescribe and the habits that behave like drugs, because a concentration can move without anyone changing the prescription. Margin asks how much distance sits between a useful concentration and a harmful one, which decides whether the drug can be managed by clinical judgement alone or needs a measured level. The single feature a stem plants to make one of these questions decisive is its **pharmacological discriminator**, and well-written stems are built around exactly one.

RULE **In a pharmacology stem, the answer is the mechanism, not the diagnosis.** If your answer to a stem in this territory can be written without naming a receptor, a half-life, an enzyme or a therapeutic margin, you have almost certainly answered the wrong question. The diagnosis in the stem is usually scenery: it tells you which class the patient is on, and nothing more. What earns the marks is the pharmacological property that makes this patient's problem inevitable and the within-class or within-schedule move that follows from it.

27.2 The vignette that resolves on a binding profile

The commonest stem in this territory describes a patient who is responding to treatment but cannot tolerate it. She is sedated all morning, or has gained weight steadily, or feels her heart racing when she stands, or has a dry mouth and constipation, or has developed galactorrhoea, or cannot sit still. Every one of those complaints is a receptor being occupied, and the reading move is **reverse inference**: take the symptom the patient reports, ask which receptor family generates it, then ask whether the prescribed agent has meaningful affinity there. The full map from binding profile to predicted adverse effect is developed in its own right within this chapter, and the biochemistry of the transmitter systems underneath it belongs to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

What makes the stem answerable rather than merely descriptive is that the receptor families cluster into recognisable symptom groups. Histaminergic and cholinergic occupancy produce one kind of complaint, adrenergic occupancy another, dopaminergic occupancy a third, and the serotonergic actions a fourth. Once the symptom has been mapped back to a family, the answer is a within-class move to an agent whose affinity at that family is lower, not a change of class and not a change of diagnosis. That is why the comparative profiles assembled elsewhere in this chapter are worth learning as a ranked shape rather than as a list of drugs: the vignette is asking you to substitute along one axis while holding the others.

TRAP **Treating members of a class as interchangeable because they share a label.** Candidates who have learned the classes as boxes will answer that any agent in the box will do, and lose the mark. Agents within a single class differ in what else they bind, how long they last and which enzymes clear them, and those differences are precisely what the stem is testing. **Binding promiscuity** is the reason two drugs with the same primary mechanism produce entirely different complaints, and the reason a within-class switch is a real therapeutic manoeuvre rather than a cosmetic one.

27.3 The vignette that resolves on a half-life

A second family of stems turns entirely on how long the drug stays. The patient misses a few doses and becomes unwell in a characteristic way; or stops the drug abruptly and develops symptoms that could be relapse or could be discontinuation; or has a level checked soon after a dose change and the clinician acts on it; or is switched from one agent to another and something goes wrong in the overlap. In every case the discriminator is the elimination half-life and the accumulation that follows from it, developed in full among this chapter's pharmacokinetic principles, with the underlying kinetic derivations in the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

Three consequences do most of the work in vignettes. A drug that is eliminated slowly buffers itself: missed doses matter less, the plasma concentration changes gradually, and symptoms on stopping are muted and delayed, so a picture arising soon after cessation is more likely to be something other than withdrawal. A drug eliminated quickly does the opposite, producing interdose symptoms, an abrupt **offset** and a withdrawal picture that arrives fast. And accumulation takes several half-lives whatever the drug, which means any judgement made before the concentration has plateaued, whether clinical or laboratory, is a judgement about a moving number. The interval a clinician must leave empty between stopping one agent and starting another, the **washout**, is the same arithmetic used in the opposite direction, and it is the reason some switches are simple and others are not.

PITFALL

Sampling a drug level before the concentration has plateaued, and then changing the dose on the result. A level drawn too early after starting or after a dose change is not a steady-state level, it is a point on a rising curve, and treating it as though it were the patient's true exposure leads to a dose increase in a patient who was already on their way to an adequate concentration. The same error in reverse explains apparently unstable levels in a patient whose adherence is intermittent. The remedy is arithmetic, not repetition: know the drug's half-life, know when the plateau is due, and sample against it.

27.4 The vignette that resolves on an enzyme

The third family of stems is the one where nothing about the psychotropic prescription changed and the patient changed anyway. A physician added an antimicrobial or an antifungal; a cardiologist started something new; the patient bought a herbal preparation over the counter; the patient stopped smoking on admission to a ward where smoking is not allowed. Then a drug that had been tolerated for months becomes toxic, or a drug that had been working stops working, and the stem waits to see whether the candidate looks past the psychiatric drug list.

The direction of reasoning is simple and it is worth having automatic. An added inhibitor reduces the clearance of a substrate, so the substrate's concentration rises, and it rises promptly, because inhibition acts at an enzyme that is already present. An added inducer increases the amount of enzyme, so the substrate's concentration falls, and it falls with a lag, because new enzyme has to be synthesised, and it recovers over a comparable but somewhat longer interval when the inducer is withdrawn, the figures for both directions being carried by this chapter's cytochrome P450 section. That asymmetry of tempo is itself a discriminator: an abrupt deterioration points toward inhibition, a loss of effect that develops over days rather than hours toward induction. The clinically important variant is **phenoconversion**, in which a patient with entirely normal metabolic capacity is turned by a co-prescribed inhibitor into the functional equivalent of a poor metaboliser. The psychotropic substrate map itself is developed within this chapter; the interaction matrix spanning several classes at once belongs to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

GOOD TO REMEMBER

When a stable patient destabilises, date the change against the whole drug list, not the psychiatric one. Ask what a non-psychiatric colleague started or stopped, what the patient bought without a prescription, and whether tobacco use changed, because admission to hospital silently removes an inducer from many patients. A symptom whose onset lines up with a change somewhere else in the list is an interaction until proved otherwise, and the question that resolves it is which enzyme both drugs share.

27.5 The vignette that resolves on the margin

Some stems are constructed so that the pharmacology is unremarkable except for one property: the drug has very little distance between the concentration that treats and the concentration that harms. Then a modest kinetic disturbance, one that would be trivial for a forgiving drug, carries the patient across the line. A **narrow therapeutic index** is what converts an ordinary intercurrent illness, a change in fluid balance, or a new co-prescription into a toxic presentation, and it is the single property that justifies routinely measuring a concentration at all. This chapter develops the index itself, and which psychotropics need a measured level, as a principle; the target concentrations, and the monitoring framework built on them, belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

The reading move is to notice that the stem has told you about the patient's physiology rather than about the psychiatry. Anything that alters volume status, reduces renal clearance or competes for the elimination route will move the concentration of a narrow-index drug while the prescription stays identical, and for a drug with a wide margin the same events pass unnoticed. Which precipitants matter for which agent is developed within this chapter's sections on the narrow-index drugs, and the answer a stem wants is rarely to stop the drug reflexively; it is to name the kinetic precipitant, state that the concentration will have moved, and act on that. Toxic presentations themselves, and the psychotropic emergencies more broadly, are recognised and managed under the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

27.6 The vignette that resolves on intrinsic activity

A subtler stem describes a patient in whom a drug was added and something got worse rather than better, or in whom two agents acting on the same system produced an effect neither would have produced alone. The property being tested is **intrinsic activity**: whether the agent at a given receptor is a full agonist, a partial agonist, a plain antagonist or an inverse agonist. A partial agonist behaves in opposite directions depending on what the receptor is already seeing, damping the system where tone is high and supporting it where tone is low, which is why adding one to a patient already maintained on a full antagonist can produce an apparent loss of effect. This chapter teaches partial agonism as a prescribing variable in full, and the receptor theory beneath it belongs to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

The same reasoning explains a second stem shape, in which an agent modulates the action of an endogenous ligand rather than blocking it, so its behaviour depends on the state of the system rather than on the dose alone. With this axis in hand, the question is not how much drug is present but what the receptor was doing before it arrived.

27.7 The vignette that resolves on tempo

When a problem starts is often more informative than what it looks like. Adverse effects have characteristic latencies, and the **latency** a stem specifies is frequently the discriminator that assigns the phenomenon to a mechanism. Something that appears within hours of a first exposure, before any accumulation has occurred, is a direct pharmacodynamic effect of occupancy. Something that emerges over the first weeks, as the concentration plateaus and adaptive changes begin, is typically dose-related and often modifiable by dose. Something that appears only after months or years of continuous exposure implies receptor adaptation rather than occupancy alone, and is the tempo of the late movement disorders whose spectrum this chapter develops in full. Something whose onset bears no relation to dose or duration at all is **idiosyncratic**, and cannot be titrated around.

Tempo also separates the two commonest look-alikes in this whole territory. Symptoms that appear within a short and predictable interval of stopping a drug, an interval read against that drug's half-life, and that resolve when the drug is reinstated, are a withdrawal or discontinuation phenomenon. Symptoms that appear later, at a remove the kinetics do not account for, and that do not remit promptly on reinstatement, are relapse. Do not separate them by what they look like. Where a drug was prescribed for the very state it suppresses, withdrawal reproduces the original presentation almost exactly, a trap the benzodiazepine section works through in detail. The discrimination is kinetic and pharmacological: the interval from cessation, and the response to reinstatement. Where the phenomenon is a substance use disorder in its own right rather than a pharmacological adaptation, the assessment and treatment belong to the Substance use disorders notes, and the clinical use of the anxiolytics and hypnotics belongs to the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes.

27.8 Look-alike stems and the reading spine

Consolidation is easiest as pairs. The table sets out the stems most often confused, against the property that separates them and the axis the reasoning lands on.

TWO STEMS THAT LOOK ALIKE	THE PHARMACOLOGICAL DISCRIMINATOR	THE AXIS IT LANDS ON
Symptoms soon after stopping a drug against symptoms emerging later, the two pictures looking much alike	Interval from cessation, and whether reinstatement resolves it	Half-life and offset
Loss of effect over weeks against abrupt toxicity in a stable patient	Whether the added agent induces or inhibits	Enzyme interaction
Intolerable sedation against intolerable restlessness on the same class	Which receptor family the agent occupies most avidly	Binding profile
An apparently unstable drug level against a genuinely rising one	Whether the sample was drawn at plateau	Accumulation kinetics
Toxicity after an intercurrent illness against toxicity after a dose increase	Whether a kinetic precipitant altered clearance	Therapeutic margin
Worsening after adding a second agent to an established one against worsening on the first alone	Full antagonism against partial agonism at the shared receptor	Intrinsic activity
An adverse effect present from the first exposure against one emerging after prolonged treatment	Latency, and its relation to dose	Tempo of onset

MNEMONIC

What it binds, How long it lasts, Alongside what else, Tolerance of error

The four questions to run across any pharmacology stem: **W**hat does the agent bind, and with what relative affinity; **H**ow long does it last, and has it reached plateau; **A**longside what else is the patient taking, prescribed or not; and what **T**olerance of error does the drug allow between a useful and a harmful concentration. They spell WHAT, which is the question the stem is really asking, and each one is a door to a different section of this chapter.

With the discriminator identified, a marks-winning answer follows four moves in order.

- Name the pharmacological property the stem has planted, in pharmacological language.
- State the mechanism that connects that property to the patient's problem.
- Choose the move that acts on that mechanism: a within-class substitution, a change of interval or schedule, a washout, or a measured concentration.

- Say explicitly what you would not do, and why the obvious algorithmic answer does not address the mechanism.

That final move is where the pharmacology becomes visible to an examiner. A candidate who says the drug should be stopped has answered a question about a patient; one who says it should not, because the problem is an inhibitor that will wash out or a withdrawal effect that reinstatement will resolve, has answered a question about pharmacology. Every vignette here rewards the second answer, worked from the four axes without a treatment algorithm.

> In a pharmacology vignette, identify the property the stem has planted - binding, duration, company or margin - and answer with the mechanism it implies, not with a treatment algorithm.

28 · Consolidation and quick review

Every class in this chapter holds several agents that share a mechanism and differ in what they cost the patient. That sentence is the chapter compressed. Each class was assembled around one shared action: blockade of the dopamine receptor, inhibition of monoamine reuptake, potentiation of the inhibitory chloride channel, inhibition of the metabolising enzyme. The shared action makes the agents members of a class and explains why they work at all. Everything else about them, what else the molecule binds, how long it stays in the body, and how much room lies between the exposure that helps and the exposure that harms, is where the members separate. A reader who knows the class knows why the drugs work and still cannot say which of them to give.

This closing section teaches no new agent, because every drug has already been taught in the section that owns it. What remains is to state the argument those sections were arranged to make, and to convert it into something that can be produced from a blank page. The argument is that **within-class selection** is a prediction rather than a preference: given a drug's binding profile, its handling by the body and the width of its safety margin, most of what separates it from its neighbours in the same class can be worked out before any of it is observed in a patient. The pharmacology needed to make that prediction is small, finite, and the same for every class here.

28.1 One argument, tested five times

The chapter moves through receptor and kinetic principles, then the antipsychotics, the antidepressants, the mood stabilisers, and the anxiolytics, hypnotics and remaining agents. Read as five territories they are a long list to be survived. Read as one argument they are a single claim tested five times, and the claim is that the difference between two members of a class is generated by pharmacology that can be stated in advance. The principles are not an introduction to the four drug bands; they are the instrument those bands are read with.

The claim has a testable shape, which is what makes it worth holding. Take two agents from one class and set out, for each, what it binds beyond the shared target, how quickly it is cleared and by what route, and how much distance lies between the dose that works and the dose that harms.

Those three statements anticipate most of what the patient reports, most of what the prescriber has to manage, and most of the practical difference in how the two drugs are started, stopped and exchanged. What they do not anticipate is which agent a unit happens to stock, or which one the reader saw used most often as a junior, and those two considerations decide a great many real prescriptions.

Each pairwise difference runs along a **separating axis**: one pharmacological property on which the members of a class can be ranked against each other. The axes are few and they recur. Sedation, weight, prolactin and movement liability rank the antipsychotics. Half-life, enzyme inhibition and binding beyond the transporter rank the antidepressants. Index width, elimination route and reproductive risk rank the mood stabilisers. Receptor subunit selectivity and half-life rank the hypnotics. Naming the axis rather than the drug converts a remembered preference into a defensible choice.

RULE **The shared mechanism explains the effect; everything that separates the members of a class is a side action, a kinetic difference or a margin.** Establish what two agents have in common only to confirm they belong to the same class. The clinical question is always what they do not share, and that question has a small and recurring set of answers.

28.2 Five questions asked of every agent

Because the separating axes recur, a single **retrieval card** serves the whole chapter. Rather than storing each agent as an unstructured block of facts, the reader learns one set of prompts and answers them for the drug in hand, which reduces the memory load and guarantees that no limb of an answer is dropped under time pressure. They spell CLASS, which is the level this chapter works at and the level at which an answer here should be pitched.

MNEMONIC

Clearance, Ligand behaviour, Adverse-effect signature, Safety margin, Selection

For any agent in any class, retrieve: **C**learance, meaning the half-life, the route of elimination and the enzymes that handle the drug; **L**igand behaviour, meaning whether the molecule acts as a full agonist, an antagonist, a partial agonist, a reuptake inhibitor or an enzyme inhibitor, and whether that action is reversible; **A**dverse-effect signature, meaning what else the molecule binds and therefore what the patient will report; **S**afety margin, meaning the distance between the exposure that treats and the exposure that harms; and **S**election, meaning the within-class choice that all four preceding answers were gathered to justify. Read together they spell CLASS, which is exactly the level this chapter teaches.

The card is built out of constructs the chapter has already supplied, so nothing in it has to be learned separately. Its value is that it disciplines an answer. A reader who has produced all five prompts for the agent in front of them has generated a structured response: how the drug is handled, what it does at the receptor, what the patient will feel, how much room for error exists,

and what follows for the choice. An examiner reads that as pharmacological reasoning rather than as recall of a list.

Binds	**Stays**	**Margin**	**Chooses**
WHAT ELSE IT TOUCHES	HOW LONG IT LASTS	ROOM BEFORE HARM	THE AGENT THAT FOLLOWS

28.3 What each question actually predicts

The card is only useful if each prompt is attached to the decision it settles, because a prompt that returns a fact without a consequence is one more thing to memorise. Set out this way, the four pharmacological questions each buy a different practical answer and the fifth is the answer itself. The table is a retrieval index rather than a teaching table: the full account of each property lives in the section that owns it.

THE PROMPT	WHAT THE ANSWER PREDICTS	WHERE MEMBERS OF A CLASS DIVERGE
Clearance: half-life, elimination route, metabolising enzyme	Dosing interval, time to a steady concentration, whether a measured level means anything yet, how abruptly the drug may be stopped, and how a switch is staged	Agents sharing a mechanism still differ on stopping and on switching, because the shorter-acting member falls away faster
Ligand behaviour: agonist, antagonist, partial agonist, reuptake inhibitor, enzyme inhibitor, and reversibility	Whether the effect has a ceiling, how the system behaves on withdrawal, and whether one molecule can look stimulating in one state and blocking in another	A partial agonist and a full antagonist at the same receptor share a class and are not interchangeable
Adverse-effect signature: what else the molecule binds	The symptoms the patient will report, anticipated before any appear, and the tolerability conversation held before the first prescription	Promiscuous binding predicts a broad adverse-effect signature; a cleaner profile predicts a narrower one at the same efficacy
Safety margin: distance between effective and toxic exposure	Whether the drug is titrated to clinical effect or to a measured concentration, and how much a small change in fluid balance or renal handling matters	A narrow margin makes an ordinary event dangerous where a wide margin absorbs it
Selection: the member that fits this patient	The actual prescription, defended by the four answers above rather than by habit	This is the output, not another property to memorise

Worked in both directions, given an agent produce its properties and given a property produce the agents at either end of it, the index becomes a network rather than a sequence.

28.4 Retrieve across the classes, not down one

A chapter learned in the order it was read is fragile, because recall then depends on starting at the top of a class and running down it. Durable retrieval comes from asking the same question of every class in turn, which is **interleaved retrieval**: deliberately mixing the classes during practice

instead of drilling one to exhaustion. It feels slower than working through a single class at a time, and that difficulty is why it works, because it trains the discrimination a stem actually tests.

Two prompts are worth interleaving across all five sub-areas. The first is the separating axis. The second is the **class-defining hazard**: the adverse consequence characteristic of the whole class rather than of one member, which therefore has to be considered whenever any member is prescribed. Together they explain both why a class is used and why choosing carelessly inside it has consequences.

CLASS	WHAT SEPARATES THE MEMBERS	THE CLASS-DEFINING HAZARD
Antipsychotics	Sedation, weight, prolactin and movement liability, each read off the binding profile rather than off the older typical and atypical labels	Movement disorder from dopamine blockade, and metabolic burden from the wider binding profile
Antidepressants	Half-life, inhibition of the metabolising enzymes, and how much the molecule binds beyond its intended transporter or enzyme	Disturbance on stopping, which tracks the half-life, and interaction risk, which tracks enzyme inhibition
Mood stabilisers	Width of the therapeutic index, route of elimination, whether the agent induces its own metabolism, and reproductive risk	Toxicity in a narrow-index agent, and harm to a developing fetus with the agents that carry it
Anxiolytics and hypnotics	Receptor subunit selectivity and half-life, which together set both the daytime residue and the rebound	Tolerance, dependence and withdrawal as pharmacological consequences of sustained receptor exposure
Stimulants and the non-stimulant agents	Speed of onset, duration of action, and what the formulation does to the concentration curve	Misuse and diversion, which follow rapidity of onset rather than the molecule alone

As a drill this means posing the class as the question and producing the axis and the hazard from memory, then reversing it: given an axis, name the class it ranks; given a hazard, name the receptor action that generates it. That movement, from a mechanism to its consequences and back, is what a viva is built to elicit.

28.5 The failure mode this chapter exists to correct

Set against the pharmacological method is the alternative most prescribing actually runs on, the **familiarity heuristic**: choosing the agent one has used before, seen used, or been taught by whoever supervised the first ward. Familiarity is not worthless, since a drug the prescriber knows well is monitored better than one they do not. It is simply not a reason, and it cannot be defended when the choice turns out badly. The remedy is not to abandon experience but to state the axis experience was tracking.

The examination version of the same failure is answering a within-class question with a name instead of a property. The stem describes a patient who cannot tolerate sedation, or who is

unlikely to attend for a blood test, or in whom a missed dose is likely; the answer that scores states the axis first and reaches the agent second. The antidote is **cold retrieval**, producing the card from a blank page with the book closed rather than reviewing an answer already written, because recognising a correct choice on the page is a far weaker skill than generating one from a bare stem.

TRAP **Answering a within-class question with a drug name instead of the property that selects it.** Candidates do this because the name is what they revised and because it feels decisive. The examiner is testing whether the choice can be derived, so an answer that names the separating axis and then lands on an agent earns marks an unsupported name does not, even when both arrive at the same drug.

PITFALL

Switching within a class and expecting the troublesome adverse effect to disappear, when the two agents share the binding that generated it. A switch is worth making only when the profiles genuinely differ on the axis causing the problem; where they do not, the patient endures a second titration and arrives at the same complaint. Before any within-class switch, state which property is expected to change and by what mechanism. If that sentence cannot be completed, the switch is being made on familiarity.

28.6 Where this chapter stops

Consolidating a chapter includes consolidating its edges, because a disciplined answer is defined as much by what it hands over as by what it teaches. The **handover boundary** matters unusually here, since the systematic pharmacology of a class sits beside several chapters that treat the same drugs from a different angle. Each of the following should be retrieved as an interface: named where it touches the question, handed to its owner, and left there.

- The underlying science of receptor theory, the kinetic equations and the enzymology of drug metabolism belongs to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes; this chapter uses those principles as instruments rather than deriving them.
- Choosing and managing an antipsychotic for a patient with a psychotic illness, including the management of movement disorder in that setting, belongs to the Schizophrenia: management, antipsychotics and adverse effects notes.
- Choosing an antidepressant or a mood stabiliser for a depressive episode or for bipolar disorder, and the clinical monitoring that goes with it, belongs to the Mood disorders: pharmacotherapy notes.
- The clinical use of anxiolytics and hypnotics belongs to the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes, while sedative dependence as a diagnosable disorder, with its assessment and detoxification, belongs to the Substance use disorders notes.
- Indication-specific dosing for attention deficit hyperactivity disorder belongs to the ADHD and specific learning/communication disorders notes; the antiepileptic lens on the

anticonvulsant mood stabilisers belongs to the Epilepsy and psychiatry notes; and the cognitive enhancers as treatments belong to the Dementias and neurocognitive disorders notes.

- Interactions spanning classes, prescribing in physical illness and in pregnancy, monitoring schedules running across psychotropics, therapeutic target ranges, and the drug emergencies all belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

Retrieved as a rehearsed set, the boundary supplies a ready-made handover sentence: name the neighbouring construct, say why it matters to the case in front of you, point to where it is taught, and return to the pharmacological question actually asked.

28.7 Weighting the comparison to the practice you will work in

Not every within-class comparison carries the same weight at the bedside you will actually stand at. Consolidation is finite, so the last allocation decision is where to spend the retrieval that remains. Weight it toward the comparisons that recur locally and the agents genuinely available to prescribe, rehearsing those pairs until the axis separating them comes without effort, and let rarely used agents sit at lower priority.

IN INDIA

Two pharmacogenetic differences taught in the sections that own them change a within-class decision on Indian ground, and both should be retrieved as decisions rather than as facts. Where a variant that slows clearance is documented as common in Asian populations, the decision is to open at the lower end of the range and titrate against response rather than to start at a standard dose and react to the adverse effect afterwards. Where a genetic test is recommended before a specific anticonvulsant is started in people of South-East Asian ancestry, the decision is to test before prescribing rather than after a reaction has occurred. Both remain within-class decisions, because in each case a differently handled member of the same class stays available.

28.8 Carrying the frame into the examination

If cold retrieval is the engine, **spaced retrieval** is the schedule that makes it last: returning to a class after a deliberate gap rather than in one continuous block. A memory retrieved just as it begins to fade is strengthened more than one retrieved while still fresh, so letting a class cool between attempts is a feature of the method rather than neglect. Across a revision period the gaps before returning to a secure class lengthen while a shaky one is brought forward, which concentrates limited hours on the gaps the reader actually has.

GOOD TO REMEMBER

Finish every within-class decision with a spoken sentence: this agent rather than that one, because of this property, at this cost. Said aloud in clinic it exposes a choice made on familiarity in a way that thinking it silently does not, and the same sentence is the one that earns marks in a viva. A decision that cannot be finished in that form has not yet been made pharmacologically.

Carried into the examination, the chapter collapses to a sequence that survives when detail fails. Identify the class from the mechanism, run the card for the agent in question, name the axis that separates the members, and state the class-defining hazard that constrains all of them. What consolidation buys is not more facts but faster and more reliable access to the facts already learned, so that under pressure the pharmacology arrives as a method rather than as a set of half-remembered associations about individual drugs.

> Within a class the mechanism is shared, so the choice is made from what differs: what else the agent binds, how fast it clears, and how much room lies between the exposure that helps and the exposure that harms.

Sources and further reading

Every dose, threshold, half-life and frequency in this chapter is grounded in one of the works below; where two sources set different figures, both are named in the text as a conflict rather than reconciled. Rating scales and monitoring instruments are described and cross-referred, never reproduced. Editions and publication years here were read from each work's own title and copyright pages rather than recalled.

Taylor DM, Barnes TRE, Young AH.
The Maudsley Prescribing Guidelines in Psychiatry. John Wiley and Sons. The 15th edition, first published 2025, is the chapter's principal prescribing source; the 14th edition, first published 2021, is retained where it carries a statement the 15th drops. The Maudsley Deprescribing Guidelines, first published 2024, is the source for withdrawal and tapering.

Sadock BJ, Sadock VA, Ruiz P.
Kaplan and Sadock's Comprehensive Textbook of Psychiatry. Eleventh edition, Wolters Kluwer, 2025. The principal reference across the chapter: receptor pharmacology, the antipsychotic and antidepressant classes, and the mood stabilisers. Its India chapter is cited separately where local prescribing practice is at issue.

Stahl SM.
Stahl's Essential Psychopharmacology: Neuroscientific Basis and Practical Applications, and the companion Prescriber's Guide, seventh edition, Cambridge University Press, 2021. Receptor mechanism and the profile-to-effect reasoning that organises this chapter.

Meyer JM, Stahl SM.
Stahl's The Clozapine Handbook, Cambridge University Press, first published 2020; and The Lithium Handbook, Cambridge University Press and Assessment, first published 2023. The two single-drug monographs behind the clozapine and lithium sections.

Brunton LL, Knollmann BC (Goodman and Gilman).
Goodman and Gilman's The Pharmacological Basis of Therapeutics, fourteenth edition. Receptor theory, dose-response and the therapeutic index as pharmacology rather than as prescribing. The publisher and year are deliberately not printed: the parsed copy held in this project has no title page and no copyright page, so neither could be read from the record, and neither is stated here from memory.

Tripathi KD.
Essentials of Medical Pharmacology. Eighth edition, Jaypee Brothers Medical Publishers, 2019. The pharmacology reference in widest use in Indian training, cited for class mechanism and for agents whose Indian use differs from the Western default.

Johnstone EC, Owens DGC, Lawrie SM, et al.
Companion to Psychiatric Studies. Eighth edition, Elsevier, Churchill Livingstone imprint, 2010. Class structure and the comparative pharmacology of the antidepressants.

Tasman A, Kay J, Lieberman JA, et al.
Tasman's Psychiatry. Fourth edition, John Wiley and Sons, first published 2015. Comparative within-class parameters, including the stimulant duration-of-action table.

Martin A, Bloch MH, Volkmar FR (Lewis).
Lewis's Child and Adolescent Psychiatry: A Comprehensive Textbook. Fifth edition, Wolters Kluwer, 2018. Cited for pharmacokinetic principle and for the stimulant pharmacology this chapter teaches at class level.

Indian Psychiatric Society.
Clinical Practice Guidelines. The source for the within-class Indian rankings of weight-gain and prolactin liability used in this chapter.

Also cited, in single passages.
Practical Psychopharmacology: Clinical Topics in Psychology and Psychiatry; the Shorter Oxford Textbook of Psychiatry; Stahl's Substance Use; and The Maudsley Prescribing Guidelines for Mental Health Conditions in Physical Illness. Each grounds one or two claims only, and each is named here by title alone because its edition and publisher were not read from the copy held in this project.

PART VIII

VIII

Psychopharmacology II: Emergencies, Resistance, Monitoring and Interactions

Prescribing when it goes wrong or will not work: the emergencies, treatment resistance, monitoring, and the interactions that cause both.

01 Drug emergencies

02 Treatment resistance and augmentation

03 Drug monitoring, pharmacovigilance and pharmacogenomics

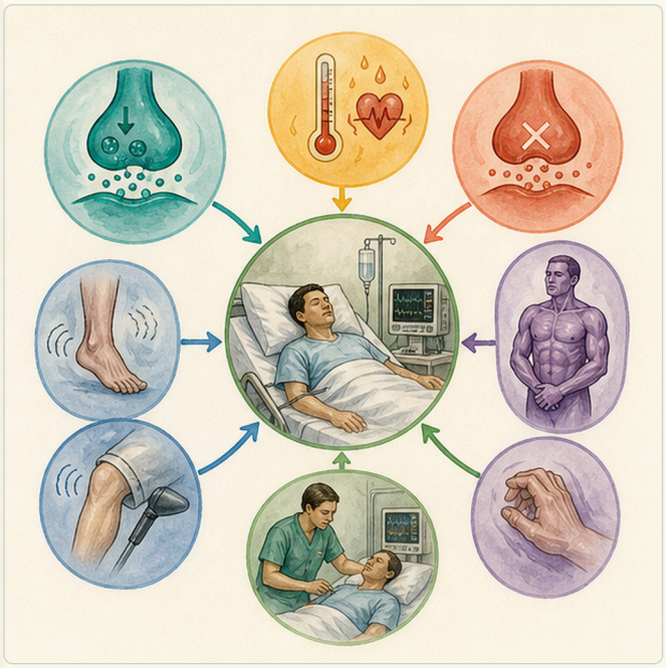
04 Interactions and prescribing in physical illness

05 Neurostimulation and device-based treatments

06 Pharmacogenetic and structural drug-reaction syndromes

07 Consolidate and apply

VISUAL 8.1



Serotonin syndrome versus neuroleptic malignant syndrome

- Serotonergic exposure
- Clonus and hyperreflexia
- Dopamine blockade
- Rigidity
- Temperature and autonomic change
- Immediate stabilisation

Serotonin syndrome and NMS are separated by exposure, onset and neuromuscular pattern while both require urgent stabilisation and culprit withdrawal.

Source lock: day-37.final.html#sec-t04

VISUAL 8.2



Treatment resistance begins with verification

- Verify
- Adequate trial
- Switch or combine
- Augment
- Rapid or neurostimulation option
- Measure outcomes

Treatment resistance is defined only after diagnosis, adherence, dose, duration and outcome have been verified, after which strategies are sequenced and measured.

Source lock: day-37.final.html#sec-t07

VISUAL 8.3



Monitoring turns risk into actionable information

- Therapeutic drug monitoring
- Blood and metabolic
- ECG and QTc
- Liver and renal
- HLA and pharmacogenomics
- Pharmacovigilance

Monitoring is useful when each test has a defined indication, timing, interpretation and action.

Source lock: day-37.final.html#sec-t12

VISUAL 8.4



Neurostimulation is a family of prescriptions

- ECT
- rTMS
- Theta-burst
- tDCS
- VNS and DBS
- Selection and monitoring

Neurostimulation modalities differ in invasiveness, target, dosing parameters, evidence and monitoring; they are prescriptions, not interchangeable devices.

Source lock: day-37.final.html#sec-t23

PAPER IV

Psychopharmacology II: emergencies, resistance, monitoring and interactions

Six bands . one complete notes document

Drug emergencies

1 · Psychopharmacology II for the psychiatrist: what this chapter owns, and what it points to

A psychotropic that was correct for the diagnosis can still be the wrong drug for the patient in front of you. The agent that suits a depressive episode meets a patient who is already taking several other medicines, whose kidneys clear the drug more slowly than any textbook assumed, whose inherited enzyme complement handles it at an unusual speed, who has already had adequate trials fail, or who returns with rigidity and fever that the prescription itself produced. In none of these has the pharmacology of the molecule changed. What has changed is that a second variable now sits between the drug and the patient, and every decision has to be made through it.

That is the whole subject of these notes. The systematic pharmacology that precedes them studies the drug class as an object in its own right: what it binds, what binding there produces, how the body handles the molecule, and what separates the members of a class from one another. This chapter takes those same molecules and asks what happens when something else is in the room. The marks here are not given for reciting a mechanism. They are given for a decision made under a named constraint, and the constraint is written into the question. Which concentration do you measure, and at what point in the dosing interval. Which agent do you avoid in this organ. Which test do you order before this prescription is written. What do you do when adequate treatment has already failed. What is this new sign, and did you cause it.

1.1 The lens: pharmacology under a second variable

What this chapter studies is the psychopharmacology of the complicated case, and its unit of teaching is the decision rather than the drug. The organising idea is **the second variable**: anything present alongside the psychotropic that changes what the correct action is. Five of them recur, and between them they account for almost every question this territory can ask. Another drug, which competes for the same clearing enzyme or adds its own effect at the same receptor. Another organ, when a heart, a liver, a kidney or a brain with an established disease alters either the exposure the patient achieves or the harm that exposure can do. Another genome, when inherited variation alters either pharmacokinetic exposure or immune susceptibility to a specific reaction. A failed trial, when adequate treatment has been given and the illness has not moved. And a syndrome the drug itself produced, when the presenting problem is no longer the disorder but the treatment.

Each of these turns a settled pharmacological fact into an unsettled clinical question. A half-life that is a property of a molecule becomes a question about this patient's clearance. A therapeutic

range that is a property of an assay becomes a question about when the sample was drawn. A familiar adverse effect becomes a question of attribution. Reading the whole chapter through that lens keeps its very different sections coherent, because the emergencies, the resistance decisions, the monitoring framework, the interaction matrix, the device-based treatments and the idiosyncratic reactions are all the same operation performed on different second variables.

RULE **No decision in this territory is made about the drug alone.** Before answering any question here, name the second variable the question has put in front of you, and say what it changes: the exposure the patient actually achieves, the harm that exposure can do in this particular organ, the susceptibility to a reaction the dose does not predict, the interpretation of a measured value, or the attribution of a new sign. A correct pharmacological statement about the molecule that ignores the second variable answers a question nobody asked.

1.2 Four operations, performed again and again

Beneath the sections lie four operations, and a reader who recognises which one a question is asking for has usually found the answer's shape before finding its content. The first is **anticipation**: predicting, before a drug is given, what a co-prescription, an organ, a genome or a pregnancy will do to it, and adjusting the choice rather than the aftermath. The second is measurement: deciding when a clinical judgement should be replaced by a number, which number, and at what moment in the dosing cycle it means anything at all. The third is **attribution**: deciding whether a new sign belongs to the illness, to a co-morbidity or to the prescription, which is the single judgement on which the management of every drug emergency in this chapter turns. The fourth is escalation: deciding what comes next when what should have worked has not, whether that is another drug, a second drug alongside the first, or a modality that is not a drug at all.

Anticipate	Measure	Attribute	Escalate
BEFORE THE PRESCRIPTION	WHEN A NUMBER BEATS A JUDGEMENT	ILLNESS OR TREATMENT	WHEN ADEQUATE TREATMENT FAILS

The operations are ordered, and the order is doing work. Anticipation is cheapest and is performed with a history and a prescription chart. Measurement is what anticipation escalates to when the variation is too large to reason about. Attribution is forced on you at the bedside, usually urgently, and is easiest for the clinician who anticipated. Escalation is the decision that follows when the first three have been done properly and the illness has still not yielded, which is precisely why a resistance judgement is only as good as the exposure and adherence work that preceded it.

1.3 The principle is taught elsewhere; the application is taught here

The first boundary this chapter keeps is with the pharmacology that precedes it. Half-life and steady state as concepts, the cytochrome system as a substrate map, the therapeutic index as a principle, and the comparison of one member of a class against another all belong to the Psychopharmacology I: principles and core drug classes notes, which teach them as properties of molecules and classes. Those notes then forward five bodies of material here by name: the target concentrations themselves, interactions that cross class boundaries, prescribing in physical illness

and across the lifespan, monitoring that spans classes, and the emergencies psychotropics produce.

The practical form of the boundary is simple. A principle is named here and moved past; an application is taught here in full. That a drug is cleared by a particular enzyme system is a principle. That adding a second agent which inhibits that system to a patient already stable on the first is an interaction with a predictable direction, a predictable time course and a defined action is an application, and applications are what this chapter is made of. The same test settles most uncertainty about where a piece of prescribing knowledge belongs: if the sentence would be true with only one drug and one healthy patient, it belongs to the class pharmacology; if it needs a second variable to be true at all, it belongs here.

TRAP **Answering a second-variable question with the class pharmacology.** Asked how to prescribe an antidepressant to a patient with established liver disease, or what to do when a mood stabiliser and a newly added analgesic meet, candidates reproduce the receptor profile and the adverse-effect signature of the class. The pharmacology may be flawless and still earn very little, because the examiner named the second variable in the stem and is marking what you did about it. State the mechanism in a clause, then spend the answer on the decision the constraint forces.

1.4 What this chapter hands over

The remaining boundaries are handovers, and each is worth rehearsing as a sentence because naming them accurately is faster and better rewarded than reconstructing a neighbouring chapter's material from memory.

- The **emergency act** is not taught here. The emergency-department presentation of a patient who is acutely unwell, the resuscitation and critical-care escalation, and the decision about where the patient is then looked after belong to the Perinatal/women's mental health and emergency psychiatry notes. This chapter owns the pharmacological lens on the same syndromes: which agents and combinations precipitate them, the mechanism by which they do, the named criteria as constructs an examiner can ask you to reproduce in your own words, the differentiation of one from another, and the drug decisions that follow.
- The **procedural act** is not taught here either. For every device-based treatment, the procedural technique is set out in the ECT and neurostimulation notes, along with the consent and legal frame and the requirements a service has to meet before it can offer the treatment at all. This chapter teaches the modality itself and its utility across disorders, at the depth a standalone answer on any single modality demands.
- Guideline-level algorithms for the treatment of a named disorder are also not this chapter's business. Where a resistance decision sits inside a disorder's own treatment sequence, the sequence belongs to the notes on that disorder, and what is taught here is the pharmacological reasoning underneath it.

Notice what the handovers have in common. In each case the neighbouring chapter owns what is done to the patient, and this chapter owns the drug reasoning that explains and follows it. That

division is not a filing convention. It is the reason an answer here should read as pharmacology rather than as emergency medicine or as a procedure list.

1.5 The map of the territory

Six territories follow, and they are arranged so that each supplies something the next assumes. The table states what each territory is for; the sections themselves supply the depth, the numbers and the attributions.

TERRITORY	THE SECOND VARIABLE IT WORKS ON	THE DECISION IT TEACHES
The drug emergencies	A syndrome the prescription itself produced	Which agents and combinations precipitate each syndrome, how to tell them apart as exam constructs, and what to stop, start and continue
Treatment resistance and augmentation	An adequate trial that has failed	What counts as resistance before the label is applied, and how switching, combining and augmenting differ as pharmacological strategies in depression and in schizophrenia
Monitoring, pharmacovigilance and pharmacogenomics	Variation you cannot see from the bedside	Which agents earn a measured concentration and when the sample is taken, how an adverse reaction becomes a recognised signal, and which inherited variation is worth testing for before prescribing
Interactions and prescribing in physical illness and across the lifespan	Another drug, another organ, another life stage	How psychotropics interact by kinetics and by pharmacodynamics, and how the choice changes in cardiac, hepatic, renal and neurological disease, in pregnancy and lactation, in the elderly and in children
Neurostimulation and device-based treatment	Illness requiring urgent definitive treatment or not responding to drugs; the entry point depends on the modality	What each modality is, how it is thought to work, which disorders it is used in and how strong the evidence is in each
Severe idiosyncratic drug reactions	An immune response nobody could predict from the dose	Which reactions are recognised, which drugs carry them, and where a pre-prescription test changes the choice of agent

Read the map in both directions. Given a patient with a complication, it says which territory holds the decision. Given a question in an examination, it says which second variable the examiner has introduced and therefore which body of reasoning is being tested. Monitoring recurs across several territories because a measured concentration, an interval on the

electrocardiogram, a blood count and a pre-prescription test are the same operation on different variables.

IN INDIA

Three of the decisions in this chapter turn on the patient rather than on the drug, and all three are worth making a routine habit at the point of prescribing rather than after a problem has appeared. Ask whether ancestry changes the pre-prescription test before certain anticonvulsants are started, and take the answer from the section on immunogenetic testing, which names the allele, the drugs and the population it applies to. Ask whether an inherited metaboliser phenotype changes the starting dose before any response has been observed, and take that from the pharmacogenomics section. And ask which agents in the prescription earn a measured concentration, and at what point relative to the dose the sample is drawn, from the therapeutic drug monitoring section. Each of the three is a question asked before the prescription is written, not a test ordered after the patient has come to harm.

1.6 Treating the prescription as a differential diagnosis

The habit this chapter is really trying to install is the habit of suspecting the treatment. A patient on psychotropic medication who develops a new physical or neurological sign has at least three candidate explanations: the psychiatric illness, an unrelated medical problem, and the prescription. Only the third is under the prescriber's direct control, and only the third can be tested by an action as simple as stopping a drug. Yet it is routinely considered last, partly because the drug has been given for weeks without trouble, partly because the patient is expected to be unwell, and partly because the medication chart is read as a record of what has been done rather than as a list of suspects.

Making the prescription a standing item on the differential changes several things at once. It brings the drug history forward from the end of the assessment to the beginning, where it is capable of altering what you do next. It makes recent changes matter more than long-standing ones, since a stable regimen is a weaker suspect than a dose increased, an agent added, or a drug stopped in the days preceding the presentation.

PITFALL

Reviewing the medication chart last. The commonest way a drug-related presentation is missed is not ignorance of the syndrome but sequence: the history, the examination and the investigations are worked through, and the prescription is checked only when nothing else has explained the picture. By then a dose has often been given again, and the account of what changed and when has been reconstructed rather than recorded. Read the chart first, note every change made in the preceding period and every agent added by another team, and ask what the combination in front of you does that neither drug does alone.

1.7 What a good answer in this territory looks like

Questions here are set as constrained problems, and they reward a recognisable shape of answer. Open by naming the second variable and saying what it changes, because that single sentence tells the examiner the question has been understood. Then give the mechanism briefly, in the

service of the decision rather than as a display. Then give the decision itself with its conditions attached, since almost nothing in this chapter is true unconditionally: a concentration means something at a stated point in the dosing interval and in a stated phase of treatment, a dose adjustment belongs to a stated degree of organ impairment, and a test belongs to a stated drug in a stated population. Finally, say what you would hand over and to whom, which demonstrates the judgement the whole chapter is built on.

GOOD TO REMEMBER

Before adding any drug to an existing psychotropic regimen, ask four questions out loud. Does the new agent share a clearing route with anything already prescribed. Does it add to an effect the existing regimen already produces, whether on the heart, the central nervous system or the same transmitter. Does the patient have an organ or a life stage that changes the exposure either drug achieves. And is there anything about this combination that should be measured rather than watched. A prescriber who answers those four before signing has performed the anticipation this chapter exists to teach, and has usually avoided the presentation the later sections are about.

1.8 A scaffold to carry into essays and vivas

Under examination pressure the method has to survive the loss of the detail. The spine that holds this chapter together is SECOND, which is a teaching scaffold rather than a validated instrument, and it will structure an answer on any topic in these notes: an emergency, a resistant illness, a measured level, an interaction, a device-based treatment or a rash.

MNEMONIC

SECOND

Second variable: name it first, whether it is another drug, an organ, a genome, a failed trial or the prescription itself. **E**xposure: ask whether the concentration in this patient is what the dose implies, and whether it can be measured. **C**ause: decide whether the new sign belongs to the illness or to the treatment before treating either. **O**ptions: when adequate treatment has failed, set out switching, combining, augmenting and the device-based modalities as alternatives with different profiles. **N**arrow margins: identify what this particular combination threatens and monitor that, rather than monitoring by habit. **D**evolve: hand the emergency act and the procedural conduct to the chapters that own them, and say so explicitly.

For a short answer, compress the scaffold to the second variable, the mechanism, the decision and the handover. For a long essay, expand each letter and draw the specifics from the section that owns them. In a viva, make the reasoning audible: say what else is on board, what that does to the exposure or to the risk, what you would therefore do differently, and at what point you would ask for help. That is the clinician the examination is looking for, and it is also the clinician who is safe with a combination no textbook has yet described.

Ask what else is present besides the drug, then say what that changes about the exposure, the interpretation, the attribution or the next step.

2 · Serotonin syndrome: precipitants, Hunter criteria, features and management

This is the emergency a prescription writes by itself. Every other syndrome in this section arrives with the illness; this one arrives with the drug chart. The drug history establishes serotonergic exposure; diagnosis then requires a compatible clinical syndrome, preferably assessed with the Hunter decision rule, and the exclusion of important mimics. The marks are correspondingly pharmacological: which agents and which combinations produce it, what the receptor account is and how firmly it may be stated, how the named criteria are constructed and what they are actually asking, and what the drug response is once the syndrome is recognised. The emergency-department presentation, the resuscitation and critical-care escalation and the disposition decision belong to the Perinatal/women's mental health and emergency psychiatry notes and are not taken up here; the comparison against neuroleptic malignant syndrome has its own section of these notes.

2.1 Why the drug chart is read first

Two properties of **serotonin toxicity** set it apart from the idiosyncratic drug reactions taught elsewhere in these notes, and both are pharmacological rather than clinical. The first is that it does not require polypharmacy. The Maudsley Prescribing Guidelines state that the syndrome can occur with a single serotonergic drug at a therapeutic dose, and more frequently with a combination of serotonergic drugs or in overdose, which means a single agent at a therapeutic dose leaves the diagnosis open. A candidate who reasons that the patient is on one antidepressant, at a licensed dose, and therefore cannot have serotonin toxicity has made the error the sentence exists to prevent.

The second is the time course. Symptoms can develop within minutes of ingestion, although most patients present within **six to 24 hours** after a medication change or overdose. That window is a useful temporal clue that must be interpreted with exposure and examination findings: it ties the presentation to a specific prescribing event rather than to the course of the illness being treated, and it tells you which page of the chart to read. An agitated, tremulous, sweating patient whose regimen changed the previous evening is a pharmacological problem until proven otherwise.

MNEMONIC

Exposure, Branch, Grade, Withdraw

The four steps in order: establish that a serotonergic agent was actually taken, find one complete qualifying criterion branch, grade the severity, and stop the offending drugs. Each step is a separate examinable construct and the sequence is not interchangeable, because the criteria set below is only licensed to run once the first step is satisfied.

2.2 The precipitants, and the routes in

The first route is additive pharmacodynamics, and one pairing dominates it. The Maudsley guidelines state that most severe cases involve an MAOI, including moclobemide, plus an SSRI. The detail that carries the mark is the parenthesis: **moclobemide** is named inside the

monoamine oxidase inhibitor group rather than outside it, so a candidate who treats it as a safer member of the class and pairs it with a reuptake inhibitor has walked into the exact combination the source singles out.

Next come selected non-psychiatric precipitants, which are not a separate mechanism at all. They are a prescriber and indication grouping, and the agents in it act through the same additive pharmacodynamics, by reuptake inhibition, monoamine oxidase inhibition or serotonin release. The American Family Physician review of serotonin syndrome tabulates the medications that may contribute, and separates an antidepressant and mood-stabiliser class from a Miscellaneous class of five entries:

- Cocaine
- Dextromethorphan
- **Linezolid**
- L-tryptophan
- 5-hydroxytryptophan

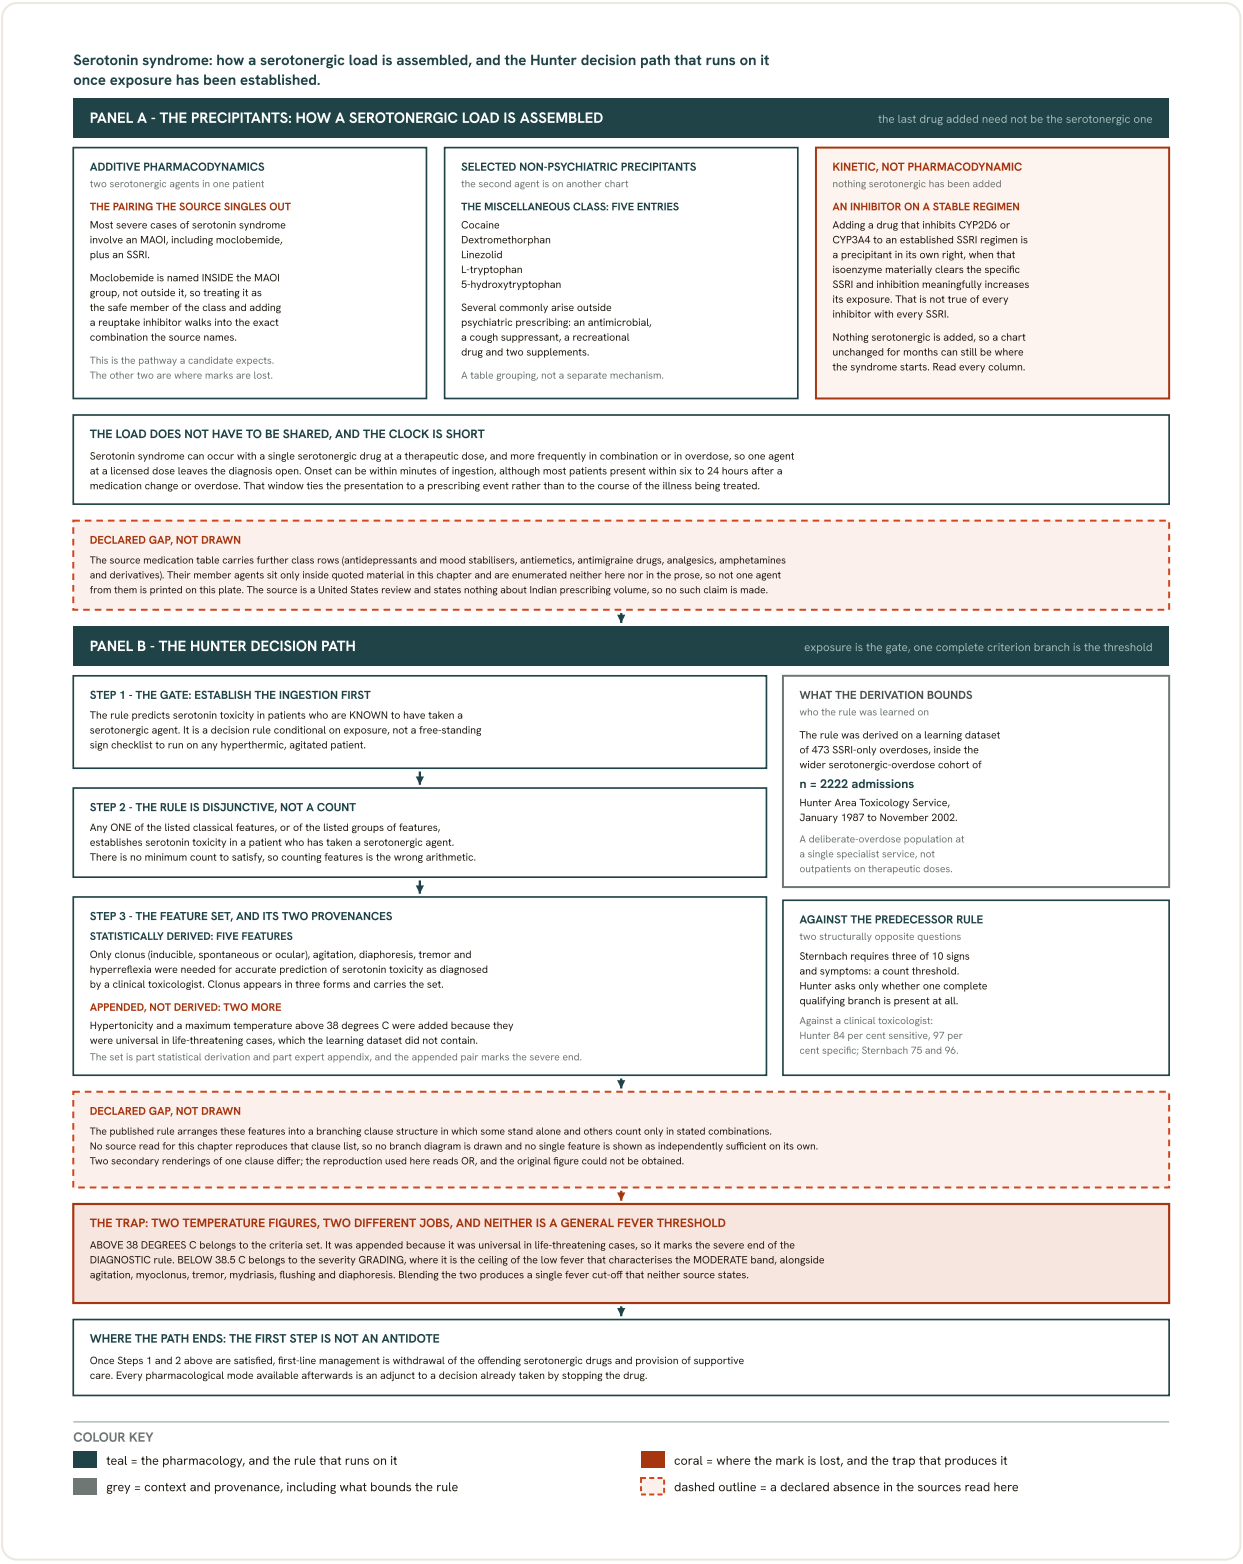


Fig 37.1 Serotonin syndrome: how a serotonergic load is assembled, and the Hunter decision path that runs on it. A single agent at a therapeutic dose leaves the diagnosis open, and most patients present within six to 24 hours of a medication change or overdose. Most severe cases involve an MAOI, including moclobemide, plus an SSRI. Selected non-psychiatric precipitants are a prescriber and indication grouping, not a separate mechanism; only the Miscellaneous class of five is enumerated. The kinetic pathway adds nothing serotonergic: a CYP2D6 or CYP3A4 inhibitor on an established SSRI regimen raises exposure to the drug already on the chart, where that isoenzyme materially clears that SSRI and inhibition meaningfully raises it. The Hunter rule runs only once ingestion is established and is disjunctive, one complete criterion branch being sufficient; it was derived on a learning dataset of 473 SSRI-only overdoses within a cohort of 2222 admissions. The branching clause list and the unenumerated class rows are declared, not drawn. (Ables and Nagubilli, American Family Physician 2010; Dunkley and colleagues, QJM 2003; the Maudsley Prescribing Guidelines in Psychiatry.)

Read that list for what it implies about where the second drug comes from. Several entries commonly arise outside psychiatric prescribing, so the complete medication and supplement history must be reviewed: an antimicrobial and a cough suppressant sit alongside a recreational drug and two over-the-counter supplements. The serotonergic contribution that completes the picture is frequently written on another team's chart, or on no chart at all.

A further route is kinetic rather than pharmacodynamic, and it is the one this chapter exists to teach. The same review names the addition of drugs that inhibit the CYP2D6 or CYP3A4 isoenzymes to an established selective serotonin reuptake inhibitor regimen as a precipitant in its own right, when that isoenzyme materially contributes to clearance of the specific SSRI and inhibition meaningfully increases its exposure. Nothing serotonergic has been added; the exposure to the existing antidepressant has simply risen. The enzyme map itself, and the systematic cross-class interaction matrix built on it, are developed in their own sections of these notes, as is the genotype that decides how much inhibition a given patient could absorb before the interaction mattered. What belongs here is the narrower point: inhibition of an isoenzyme that materially clears the antidepressant already on the chart can itself be the precipitant, so the last drug added is not necessarily the serotonergic one, and the underlying principles of enzyme kinetics and steady state belong to the Psychopharmacology I: principles and core drug classes notes.

> **GOOD TO REMEMBER**
>
> When the picture fits and the antidepressant has been unchanged for months, do not stop at the psychotropic column. Ask what was started in the last week for anything at all, including by another department, and ask specifically whether it inhibits either of the two isoenzymes named above. A drug chart read only for serotonergic agents will miss a possible pharmacokinetic route into this syndrome.

2.3 The receptor account, and how firmly to state it

An Indian case series of mild serotonin syndrome puts the mechanism in a single sentence: it is suggested that excess stimulation of brainstem and spinal cord 5-HT1A and 5-HT2A receptors causes the syndrome. Two things about that sentence are examinable. The first is its content, which is a receptor account and not an anatomical prediction: what is proposed is excess stimulation of those two receptor subtypes, and no more than that. The named locus belongs to the proposal itself; it is not a map from which the shape of the clinical presentation can be predicted.

The second is its strength. The source writes that it is suggested, not that it is established, and it offers this as a proposal rather than as a finding of its own series. An answer that hardens the suggestion into the mechanism has said more than the literature does, and in a viva the correct move is to give the account and then name it as proposed. Receptor pharmacology at the level of subtype selectivity and the adverse-effect signature it generates is the province of the Psychopharmacology I: principles and core drug classes notes; what this chapter owns is the toxic state that excess agonism produces when two prescriptions meet.

2.4 The Hunter criteria as an exam construct

The **Hunter Serotonin Toxicity Criteria** are examined more often for their structure than for their content, and the structure has two properties that a candidate must state explicitly. The first is a prerequisite: the rule runs only where recent serotonergic ingestion is already established. It is a decision rule conditional on exposure, not a free-standing sign checklist that can be applied to any hyperthermic, agitated patient. The second is that it is disjunctive. Any one of the listed classical features, or of the listed groups of features, establishes serotonin toxicity in a patient who has taken a serotonergic agent, and there is no minimum count to satisfy.

RULE **Exposure is the gate, and one complete criterion branch is the threshold.** Establish the serotonergic ingestion first, because the criteria are not licensed without it. Then a single complete qualifying branch is sufficient, and a branch may be one feature or a stated combination of features rather than a feature on its own. A candidate who counts features and reports that the patient has only two has applied the wrong arithmetic to the right rule.

The content came out of a statistical derivation. Only five clinical features were needed for accurate prediction of serotonin toxicity as diagnosed by a clinical toxicologist: **clonus** in its inducible, spontaneous and ocular forms, agitation, **diaphoresis**, tremor and **hyperreflexia**. Clonus is the load-bearing sign of the set, appearing in three separate forms, and it is a particularly important examination finding that appears in several branches of the published rule.

Two further features were appended rather than derived. **Hypertonicity** and a maximum temperature **above 38 degrees C** were added to the criteria because they were universal in life-threatening cases. Such patients lay outside the learning dataset. That provenance is worth carrying into an answer: the rule is part statistical and part expert appendix, and the appended pair is precisely the pair that marks the severe end. The published rule arranges these features into a branching structure in which some stand alone and others count only in stated combinations; the branch list is not reproduced here, and a candidate should not assume that every one of the five features is independently sufficient on its own.

2.5 Hunter against Sternbach, and what the derivation bounds

The predecessor rule is the comparison an examiner reaches for. The **Sternbach criteria** are a count threshold, requiring **three out of a possible 10** signs and symptoms, so the two rules are structurally opposite: one asks how many features are present, the other asks whether any qualifying feature is present at all. Against a clinical toxicologist's diagnosis the newer rule performed better on both axes at once, which is unusual, and Hunter is generally preferred because it was simpler and performed modestly better in the derivation study.

84%	97%	75%	96%
HUNTER SENSITIVITY	HUNTER SPECIFICITY	STERNBACH SENSITIVITY	STERNBACH SPECIFICITY

Where the rule came from bounds what it can be asked to do. It came out of a single toxicology service's serotonergic-overdose admissions, 2222 admissions to the Hunter Area Toxicology Service between January 1987 and November 2002, and the rule itself was derived on a learning

dataset of 473 SSRI-only overdoses inside that cohort rather than on the whole of it. Two limits follow. The population was one of deliberate overdose rather than of outpatients taking therapeutic doses, and the setting was a single specialist service rather than a general clinical population. Neither limit invalidates the rule, and both are the honest caveat to attach when a stem describes a patient who is not in either category, such as an outpatient on a stable regimen who has just been given an antimicrobial.

2.6 Grading the severity, and the two temperature figures

Severity is graded in three bands, and the grading matters pharmacologically because the treatment decision is taken against the band rather than against the criteria. The Maudsley guidelines place insomnia, anxiety, nausea, diarrhoea, hypertension, tachycardia and hyperreflexia in the mild band; agitation, myoclonus, tremor, mydriasis, flushing, diaphoresis and low fever below 38.5 C in the moderate band; and severe hyperthermia, confusion, rigidity, respiratory failure, coma and death in the severe band. Read down the three bands and the autonomic and neuromuscular features recur, more marked at each step: increasing severity is associated with more marked autonomic and neuromuscular abnormalities, including hyperthermia and rigidity. The bands describe what predominates at each level of severity, not one finding turning into another, and the features may coexist.

TRAP **Treating the two temperature figures as one threshold.** They come from different sources and do different jobs. The figure appended to the criteria set marks the temperature that was universal in life-threatening cases, so it belongs to the diagnostic rule at the severe end. The figure inside the severity grading is the ceiling of the low fever that characterises the moderate band. A candidate who blends them produces a single fever cut-off that neither source states, and then applies it to whichever question is asked.

2.7 Management, and why withdrawal is the whole of the first step

The treatment target by phase is straightforward to state and easy to under-weight in an answer. The American Family Physician review puts two things first and nothing before them: withdrawal of the offending serotonergic agents, together with supportive care. Every pharmacological mode available afterwards is an adjunct to a decision that has already been taken by stopping the drug, and the examinable point is that no antidote question arises until the offending agent is off the chart. Where the syndrome is severe, the setting of care, the resuscitation and the escalation decision belong to the Perinatal/women's mental health and emergency psychiatry notes, and this chapter stops at the prescription.

How long that first step takes, and whether it is enough by itself, is answered differently by the two sources. The American Family Physician review states that mild cases generally resolve within **24 to 72 hours** with conservative therapy and removal of the causative drugs, and the qualifier in that sentence is generally, not always. An Indian case series of mild serotonin syndrome found withdrawal alone insufficient in most cases: in most of its patients, withdrawal of the offending agents alone produced no or minimal response and cyproheptadine was then required. That does not contradict the broader American Family Physician formulation, which is withdrawal together with conservative therapy rather than withdrawal on its own, and the two

statements are printed here side by side without averaging. The split of that series between the patients who responded to withdrawal alone and those who did not is not available here and is deliberately not stated.

PITFALL

Reassuring a patient or a referring team that a mild case will settle on its own within a stated number of hours. The Western figure is qualified as a general expectation and the one Indian series in this material found withdrawal alone insufficient in most of its patients. Withdraw the drug, then review the patient against the clock rather than discharging against it.

2.8 Cyproheptadine: one drug, four published regimens, four populations

The specific agent is **cyproheptadine**, and the honest opening line about it is the one the American Family Physician review supplies: it is widely used, but there is no definitive evidence of its effectiveness in serotonin syndrome. That is not a reason to omit it from an answer, but it fixes the register. This is an adjunct with a plausible receptor rationale and weak evidence, given after the offending drug has been stopped, and an answer that presents it as the antidote of choice has overstated the literature.

For an Indian candidate the guideline figure comes first. The Indian Psychiatric Society clinical practice guideline gives **cyproheptadine 8 to 12 mg/day orally in severe serotonin syndrome**, and the same guideline's management instruction for the syndrome opens with stopping the offending drugs. The Western schedule is then named as the alternative rather than as the standard. Set the published figures side by side with the population each belongs to.

SOURCE	POPULATION THE FIGURE WAS PUBLISHED FOR	CYPROHEPTADINE AS PUBLISHED
Indian Psychiatric Society clinical practice guideline	Severe serotonin syndrome	8 to 12 mg/day orally in severe serotonin syndrome
Indian case series of mild serotonin syndrome	Mild serotonin syndrome, outpatients	8 mg three times daily, that is 24 mg/day, in mild serotonin syndrome, recorded as the amount these patients received and carried here as case-series evidence rather than as a dose the series recommends
American Family Physician review	Moderate to severe serotonin syndrome	In moderate to severe serotonin syndrome an initial 12 mg should be considered, then an additional 2 mg every two hours if symptoms continue; after the patient is stabilised 8 mg every six hours may be administered as the maintenance dosage
Indian product label for the marketed tablet	Licensed antihistaminic and antiserotonin use, the indication recorded on the label; the label carries no serotonin-syndrome indication of its own	For that labelled antihistaminic and antiserotonin indication, and for no other, the adult total daily dose is not to exceed 0.5 mg/kg/day, with a therapeutic range of 4 to 20 mg a day and the majority of patients requiring 12 to 16 mg a day. This is the marketed product's own labelled ceiling, not a serotonin-syndrome dose

Two features of that table are load-bearing. The first is the modals, which are permissive throughout the Western schedule: a loading dose should be considered and a maintenance dose may be administered, neither of which is an instruction. The second is that the product label is a label. It binds the dose of the marketed tablet and carries no serotonin-syndrome indication, so it is not the governing authority for an emergency regimen even though it is the only Indian document that states a weight-based ceiling.

A third feature is what these documents do not contain, and it is worth naming because a printed dose with nothing said above it reads like a complete instruction. The Indian Psychiatric Society entry gives a daily amount of cyproheptadine for severe cases and names neither a dose to begin with nor a ceiling above it. The American Family Physician cyproheptadine schedule states a loading dose, a titration step and a maintenance dose and then stops, so no maximum for the day appears anywhere in it. The cyproheptadine product label runs the other way: it states a range and a weight-based cap and no starting dose at all, which is what a document written to constrain a dose rather than to begin one looks like. Where a source is silent about an upper limit, that silence is the record, and this chapter does not fill it from elsewhere.

TRAP **Reading the four regimens as an ascending dose ladder.** Sorted by milligrams alone, the mildest illness appears to receive the most drug and the severe Indian figure appears to be an under-dose. That ordering is an artefact of crossing four populations, not a disagreement about severity. Each figure is only interpretable with its population attached, and an answer that lines them up as a titration schedule has manufactured a claim none of the four sources makes.

IN INDIA

Print the Indian Psychiatric Society figure first and treat the Western loading schedule as the named alternative. The marketed tablet's prescribing information does state a weight-based adult daily cap, but it states it for the licensed antihistaminic and antiserotonin indication and carries no serotonin-syndrome indication at all, so it is contextual product information rather than the ceiling that governs an emergency regimen, and dosing that runs up against it is a question for toxicology advice rather than for the label. The two documents disagree about the maximum a day may contain, and this chapter holds only the ceiling and the therapeutic range from the label, not the label's stated exceptional maximum, so the disagreement is named rather than quantified on both sides.

Exposure plus one complete qualifying criterion branch makes the diagnosis; stopping the offending drug is the treatment, and every agent given afterwards is an adjunct with its population attached.

3 · Neuroleptic malignant syndrome: features, pathophysiology and management

Neuroleptic malignant syndrome is a prescribing event before it is a critical-care event. The pharmacology question asks four things of you: which drug produced it, by what mechanism, which named criteria set the examiner has in mind, and what happens to the drug chart afterwards. The emergency-department presentation, the resuscitation and the disposition decision are taught in the Perinatal/women's mental health and emergency psychiatry notes and are deliberately left there. What follows is the drug's half of the problem, which is the half a psychiatrist owns on the ward round and the half the written paper asks about.

3.1 The prescription that produces it

The drug history is most of the diagnosis. DSM-5-TR records that affected individuals have generally met a **dopamine antagonist** within 72 hours of symptoms developing, and it takes care to present this as descriptive text rather than as a diagnostic criterion. That distinction is worth holding: read as a criterion it becomes a false reassurance whenever the exposure is older or less certain, and read as a description it simply tells you where to look first.

The offending agent is not always on the psychiatric chart. The AIIMS text Clinical Methods in Psychiatry states that the syndrome is seen more with high-potency antipsychotics and names metoclopramide and domperidone among the other offending agents. Central dopamine D2 antagonists can precipitate the syndrome whatever specialty prescribed them, and abrupt withdrawal or reduction of dopaminergic therapy is the other established route; in a rigid febrile patient the antiemetic chart is part of the drug history.

The Maudsley Prescribing Guidelines gather the risk factors into one undifferentiated row. Separating out its drug-related items is worth doing, because they describe the prescription rather than the patient:

- high-potency first-generation antipsychotics
- a recent or rapid dose increase
- a rapid dose reduction
- abrupt withdrawal of a co-prescribed anticholinergic agent
- **antipsychotic polypharmacy**

RULE **Dose change in either direction is a risk, not dose escalation alone.** The same list carries rapid increase and rapid reduction, which is hard to reconcile with a simple dose-toxicity model. These are associations the guideline reports rather than a mechanism it explains: a patient stable for years can present during a cross-taper, and withdrawal of an anticholinergic, a drug doing nothing dopaminergic, earns a place on the risk list without any account of why.

3.2 Two hypotheses, and the two drugs that follow from them

The central account begins where dopamine does thermoregulatory work. Tse and colleagues locate that role in the thermoregulatory centre of the hypothalamus, and particularly in the **anterior pre-optic nucleus**, which gives the hyperthermia a mechanism rather than leaving it as an association. The motor sign supplies a second clue. DSM-5-TR describes generalised rigidity, **lead-pipe** in its most severe form and usually unresponsive to antiparkinsonian agents. That unresponsiveness is the teaching point: this is not ordinary drug-induced parkinsonism scaled up, and an anticholinergic given in that belief costs time.

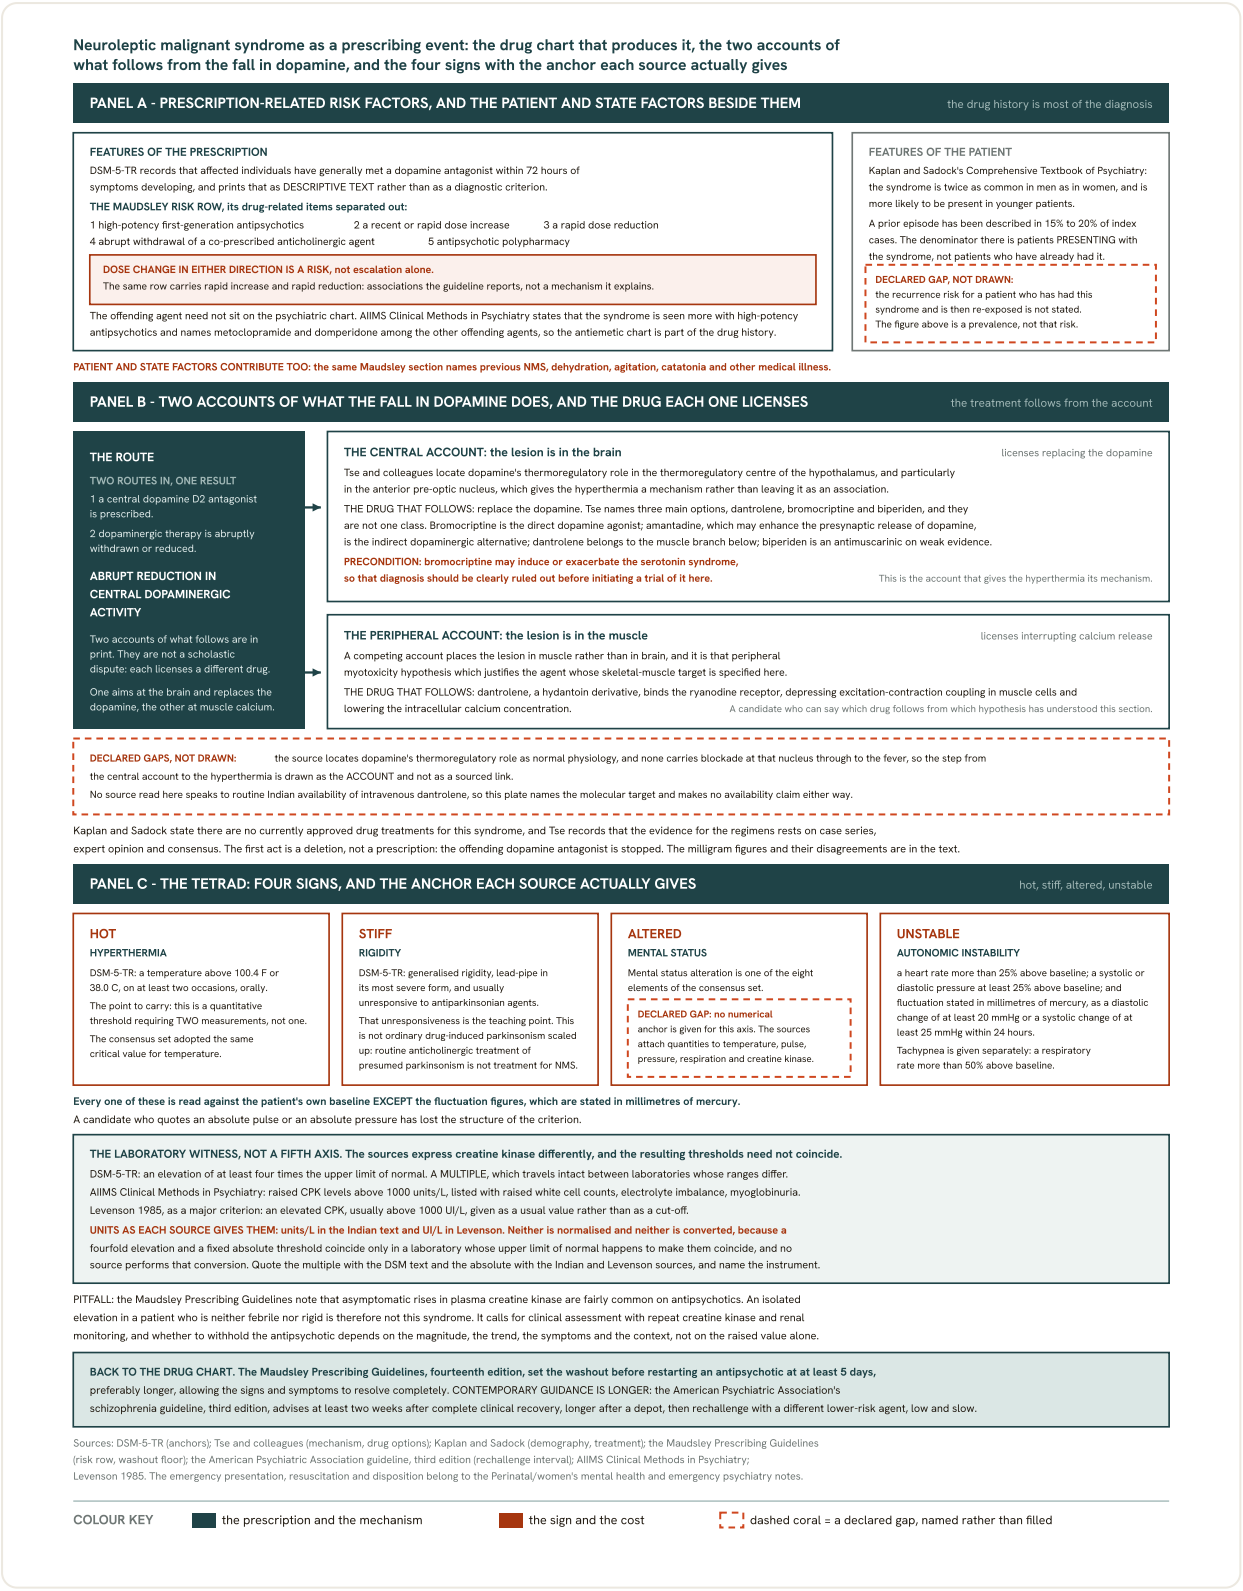


Fig 37.2 **Neuroleptic malignant syndrome: the prescription, the two accounts, and the tetrad.** Prescription-related risk factors: high-potency first-generation antipsychotics, a rapid dose change in either direction, abrupt withdrawal of a co-prescribed anticholinergic and antipsychotic polypharmacy, with patient and state factors beside them; the dopamine antagonist is generally met within 72 hours. Central dopaminergic activity falls by either route, an antagonist prescribed or dopaminergic therapy withdrawn, and each account licenses a different drug: the anterior pre-optic nucleus licenses replacing the dopamine, the myotoxicity account licenses dantrolene at the ryanodine receptor. The two creatine kinase forms need not pick out the same number: at least four times the upper limit of normal in DSM-5-TR, above 1000 units/L in the Indian text, above 1000 UI/L in Levenson. (DSM-5-TR; Tse; Kaplan and Sadock; Maudsley; AIIMS Clinical Methods; Levenson 1985.)

A competing account places the lesion in muscle rather than in brain, and it is that peripheral myotoxicity hypothesis which justifies the agent whose skeletal-muscle target is specified here. Dantrolene is a hydantoin derivative that binds the **ryanodine receptor**, depressing excitation-contraction coupling in muscle cells and lowering the intracellular calcium concentration. The two hypotheses are not a scholastic dispute, because each licenses a different drug: the central account licenses replacing dopamine, the peripheral account licenses interrupting calcium release. A candidate who can say which drug follows from which hypothesis has understood this section, and the accompanying figure carries the same route from receptor blockade to bedside sign.

3.3 Tempo: when it appears, and when it settles

Onset is early and its spread is wide. DSM-5-TR gives onset varying from hours to days after the drug is started, with some cases inside 24 hours, most in the first week and virtually all within 30 days. Recovery has a comparable shape once the drug is withdrawn: mean recovery time runs 7 to 10 days, most patients recovering within 1 week and nearly all within 30 days, the same outer bound the onset window carries.

Exposure before onset differs by drug class, though the difference is softer than the figures look. In a review of spontaneously reported cases, median exposure before onset was 23 days for cases associated with atypical antipsychotics against 6 days for typical-associated cases. Two qualifications travel with those medians and neither may be dropped. They come from a spontaneous-reporting database rather than a cohort, so they describe what was reported rather than what occurred; and the same analysis found considerable overlap between the groups, because a large share of the atypical-associated cases were taking more than one antipsychotic, frequently including a typical agent, with the reverse holding for the typical-associated cases. A clean between-class contrast is not what these data support.

3.4 How often, and in whom

Incidence depends entirely on how the cases were counted. Database studies give 0.01% to 0.02% among patients treated with antipsychotics, of the order of one case in five to ten thousand; a population-based study from Hong Kong found a higher incidence risk, and drug-safety programme data covering 1993 to 2015, reported in the Maudsley Prescribing Guidelines, give a higher figure again. These are not three competing estimates of one quantity. They are three surveillance methods, and the spread between them is the honest answer to how common the syndrome is.

0.11%	0.16%	15% to 20%	10% to 20%
HONG KONG POPULATION STUDY	DRUG-SAFETY PROGRAMME DATA	PRIOR EPISODE, INDEX CASES	FATALITY IF UNRECOGNISED

Kaplan and Sadock's Comprehensive Textbook of Psychiatry adds the demographic skew: the syndrome is twice as common in men as in women and is more likely to be present in younger patients. The mortality figure is worth little without its scope attached. DSM-5-TR puts fatality at 10% to 20% where the disorder is not recognised, which is a statement about missed diagnosis

rather than a case-fatality rate for the syndrome as a whole. Other fatality figures are in print from different eras and different case definitions; they belong to their own populations and must never be pooled into a single range.

TRAP **Reading the prior-episode figure as a recurrence risk.** A prior episode has been described in 15% to 20% of index cases, and the denominator there is patients presenting with the syndrome, not patients who have already had it. It answers how often a presenting case has a history, and says nothing about the chance that a survivor will relapse on re-exposure. That second number is the one a rechallenge conversation actually needs, and this chapter does not hold it.

3.5 The criteria sets as an exam construct

Three named sets circulate and they are built differently, which is more examinable than any single threshold. DSM-5-TR prints no counted checklist; it describes the syndrome and attaches quantitative anchors to individual features. The International Expert Consensus criteria, derived by a Delphi panel, form a counted and weighted set of **eight elements**: one drug element, recent dopamine antagonist exposure or dopamine agonist withdrawal; two cardinal signs, hyperthermia and rigidity; one cognitive element, mental status alteration; one laboratory element, creatine kinase elevation; two autonomic elements, being **sympathetic nervous system lability** and tachycardia with tachypnea counted together; and one exclusion element, a negative work-up for other causes. Levenson's 1985 criteria are older and simpler, a major-and-minor structure satisfied by all three major criteria, or two major with four minor.

CRITERIA SET	HOW DRUG EXPOSURE ENTERS	STRUCTURE	WHERE THE POSITIVE CALL IS SET
DSM-5-TR	Dopamine antagonist met within 72 hours of onset, written as descriptive text	Narrative description, with anchors attached to single features	No count and no score; judgement on the whole picture
International Expert Consensus (Delphi)	One of the eight elements, and it admits dopamine agonist withdrawal as well as antagonist exposure	Eight elements, weighted against one another by the panel	Best agreement where the total reaches 74, against modified DSM-IV-TR criteria
Levenson 1985	Absent; no drug-exposure element among either tier	Three major criteria and six minor criteria	All three major, or two major with four minor

The anchors DSM-5-TR attaches to individual features are the ones most often asked for. Hyperthermia is defined as a temperature **above 100.4 F or 38.0 C on at least two occasions, measured orally**, with profuse diaphoresis as its distinguishing companion. The autonomic criteria are partly relative and partly absolute, which is easy to garble: a heart rate more than 25% above baseline, a systolic or diastolic pressure at least 25% above baseline, but fluctuation stated in millimetres of mercury, as a diastolic change of at least 20 mmHg or a systolic change of at least 25 mmHg within 24 hours. Tachypnea is given separately, as a respiratory rate more than 50% above baseline. Every one of these is read against the patient's own baseline except the

fluctuation figures, so a candidate who quotes an absolute pulse or pressure has lost the structure of the criterion. The consensus panel's own critical values match the temperature, creatine kinase and blood-pressure-elevation anchors given here; whether it took the pulse, respiratory and fluctuation figures across as well is not verified and is not claimed.

Two further points repay attention. Levenson's set names no drug-exposure element at all, among either its major or its minor criteria, while both later sets do, so a stem offering a florid picture with no drug history is separating the sets rather than testing the syndrome. And the consensus criteria carry their own operating characteristics, sensitivity 69.6% against specificity 90.7%. Those figures must not be ranked against the decision rule taught for the serotonin syndrome in the neighbouring section: that instrument was validated in a different disease against a different reference standard, in overdose admissions rather than in archived consultations. The panel also weighted its eight elements against one another, but those per-element priority weights are not reported in the account read here and are not reconstructed from memory; the validated total is available, the weights behind it are not.

MNEMONIC

Hot, Stiff, Altered, Unstable

The four axes every set is built on: hyperthermia, rigidity, altered mental status and autonomic instability, with the creatine kinase as the laboratory witness rather than a fifth axis. They share the core clinical syndrome but differ both in the elements included and in how those elements are combined or weighted.

3.6 The creatine kinase, where the sources genuinely disagree

The sharpest disagreement in this topic is about the form of the threshold, and the two forms need not pick out the same number. DSM-5-TR sets the threshold as a multiple, an elevation of **at least four times the upper limit of normal**, which travels intact between laboratories whose reference ranges differ. The AIIMS text Clinical Methods in Psychiatry gives an absolute figure instead, raised CPK levels above 1000 units/L, listed alongside raised white cell counts, electrolyte imbalance and myoglobinuria; Levenson's major criterion is expressed the same way, an elevated CPK usually above 1000 UI/L. Both forms are in print. They are different instruments, and whether they pick out the same number depends on the laboratory's upper limit of normal; Levenson's figure is given as a usual value rather than a cut-off, and neither form is a correction of the other.

TRAP **Converting the multiple into the absolute, or the absolute into the multiple.** A fourfold elevation and a fixed absolute threshold coincide only in a laboratory whose upper limit of normal happens to make them coincide, and no source performs that conversion. Quote the multiple with the DSM text and the absolute with the Indian and Levenson sources, and say which instrument you are using. An Indian examiner is the more likely to have the absolute form in mind, which is a reason to carry both rather than a reason to substitute one.

PITFALL

Treating a raised creatine kinase as the diagnosis. The Maudsley Prescribing Guidelines note that asymptomatic rises in plasma creatine kinase are fairly common on antipsychotics. An isolated elevation in a patient who is neither febrile nor rigid is therefore not this syndrome. It requires clinical assessment with repeat creatine kinase and renal monitoring, and whether to withhold the antipsychotic depends on severity and context rather than on the raised value alone. The elevation is a supporting feature inside a clinical picture rather than a trigger standing alone.

3.7 Pharmacological management, and the rechallenge decision

Start from the honest position. Kaplan and Sadock's Comprehensive Textbook of Psychiatry states that there are no currently approved drug treatments for this syndrome, and Tse and colleagues record that the evidence supporting the different regimens rests on case series and expert opinion and consensus. Everything that follows is therefore off-label and weakly evidenced, which is precisely why the first act is a deletion rather than a prescription: the offending dopamine antagonist is stopped, and no specific agent substitutes for that. Tse names the three main drug options as **dantrolene**, **bromocriptine** and **biperiden**.

Dantrolene follows the peripheral hypothesis. For severe NMS, Kaplan and Sadock state that dantrolene may be effective, and begin with intravenous **0.8 to 2.5 mg/kg every 6 hours**, with a maximum of 10 mg/kg daily in the textbook's own words; the permissive modal is the source's own and should survive into your answer. Tse's review states the dantrolene dose in NMS differently, as intravenous 1 to 10 mg/kg body weight with no dosing interval attached, or an oral option, 50 to 600 mg once daily. These two statements do not reconcile into a single regimen and are better quoted with their sources than blended. Note also where each of them stops. The textbook's dantrolene schedule carries its own daily cap. The review's dantrolene figure states an amount and a route with no ceiling above them, so the cap printed beside the first regimen belongs to that regimen and does not travel across to the second. Neither dantrolene passage gives a separate initiation dose either: in both, the dose stated is the dose given.

Bromocriptine follows the central hypothesis, and here the same textbook contradicts itself. For NMS one chapter starts bromocriptine at 2.5 mg two or three times daily, with 45 mg as the maximum total daily dose, while another chapter of the same book carries a second, inconsistent bromocriptine regimen, added after dantrolene. That second regimen's figures are not quoted in the source behind this sentence and are not printed here; the examinable point is the contradiction itself, not the rival numbers. An independent review supports the titration form, so learn the titration and treat the disagreement as internal to the textbook rather than as a second consensus. **Amantadine**, which may enhance the presynaptic release of dopamine, is the alternative dopaminergic agent in NMS at 200 to 400 mg/day in divided doses. The passage states that range and stops there, with no initiation dose below it and no maximum above it, so the upper figure is the top of a stated range and not a stated limit.

RULE **Exclude the serotonin syndrome before giving bromocriptine.** Kaplan and Sadock state that bromocriptine may induce or exacerbate serotonin syndrome, so that diagnosis should be clearly ruled out before initiating a trial of it here. This is the clinical pay-off of the differentiation taught in the neighbouring section: the two syndromes are separated not for tidiness but because a treatment for this one worsens the other.

IN INDIA

Two decisions here are specifically Indian. First, because the AIIMS text names metoclopramide and domperidone among the offending agents and both are prescribed well outside psychiatry, the drug history in a rigid febrile patient must reach past the psychotropic chart to the antiemetics and prokinetics before any antipsychotic is exonerated or any single agent is blamed. Second, the Indian Psychiatric Society schizophrenia guideline names bromocriptine to be considered in case of NMS, inside the catatonia row of its special-conditions table, alongside intensive care admission if the clinical condition worsens and the observation that electroconvulsive therapy may be required. The table names bromocriptine for NMS but gives no NMS recommendation either for or against dantrolene, and that silence is not a recommendation against the drug. The escalation and disposition the row points at belong to the Perinatal/women's mental health and emergency psychiatry notes.

What the drug history tells you afterwards is the last piece, and it is the piece the psychopharmacology paper is most likely to ask. The fourteenth edition of the Maudsley Prescribing Guidelines sets the washout before restarting an antipsychotic at **at least 5 days, preferably longer**, allowing time for the signs and symptoms to resolve completely before **rechallenge**. Read the direction of that instruction carefully: it is a floor with an open ceiling, not a scheduled restart, and five days is not a rule that can be carried on its own. Contemporary guidance is longer. The American Psychiatric Association's schizophrenia practice guideline, third edition, advises waiting at least two weeks after complete clinical recovery, longer again where the exposure was to a depot, and then rechallenging with a different lower-risk agent at low dose and slow titration. Quote the Maudsley figure as the fourteenth edition's wording and take the two-week interval as the clinical rule.

GOOD TO REMEMBER

Two questions belong in every antipsychotic drug history and both come out of this topic. Has this patient ever had an episode of this syndrome, and how fast was the last dose change in either direction? The first is asked because a prior episode is common enough among presenting cases to be worth hunting for; the second because speed of change, up or down, sits on the risk list beside the drug itself. Neither question needs a laboratory, and both are usually answerable before the first blood sample is drawn.

Stop the dopamine antagonist first, then decide which hypothesis you are treating: dantrolene for the muscle, bromocriptine or amantadine for the dopamine, and do not give bromocriptine until the serotonin syndrome has been excluded.

4 · Serotonin syndrome versus NMS: comparative differentiation

The comparison is the question. A patient on psychotropic medication becomes febrile, altered in mental state and abnormally stiff, and two named syndromes sit on the list. Neither is excluded by how unwell the patient looks. An answer that recites the features of one syndrome and then the features of the other has not answered anything, because the examiner has not asked what each syndrome is; the examiner has asked which one this is. The marks are carried by the axes on which the two genuinely separate, by the order in which those axes are interrogated, and by an honest statement of where the separation runs out. Each syndrome is taught as an entity in its own right elsewhere in these notes. What follows is the discrimination itself, read from the pharmacology rather than from the resuscitation.

The boundary is worth fixing before anything else. The emergency-department presentation, the resuscitative and critical-care escalation and the disposition decision belong to the Perinatal/women's mental health and emergency psychiatry notes. The material here is the drug that produced the syndrome, the examination that separates the two, the criteria question as an examination construct, and what the label commits the prescriber to once it is applied.

4.1 Why this differential is not academic

The usual reason a comparison is examined is that the two entities are confusable. Here there is a second and better reason. A review of the two syndromes in Current Neuropharmacology states that there is no role for dantrolene, biperiden or bromocriptine in the therapeutic management of serotonin syndrome, and, in the same breath, no role for **cyproheptadine** in neuroleptic malignant syndrome. The syndrome-directed agents named do not cross over. Read what that does and does not say. Withdrawal of the offending drug, supportive care and the sedation of agitation are common ground, so the label does not decide the whole of management. What it decides is the syndrome-directed drug that follows, and that is enough to convert a taxonomic question into a prescribing one, because the agent chosen on the strength of the label has nothing to offer the other diagnosis if the label is wrong.

RULE **The discrimination is made from the drug chart and the neuromuscular examination, in that order.** Whatever the two presentations have in common cannot separate them, however severe it looks, so the answer is never built on how unwell the patient is. What separates them is the pharmacological class the patient was exposed to, the interval between that exposure and the illness, and the direction in which the neuromuscular system has moved. These three are the principal discriminators. Other findings may support the distinction, but none of them is individually definitive, and nothing the two presentations hold in common can do the separating at all.

Note what this rule does not claim. It does not say that the shared features are unimportant, only that a feature both syndromes produce cannot be the thing that tells them apart, and it does not promise that the three discriminating axes will always be available. In a patient brought in unaccompanied with no drug history, the class axis and the tempo axis are unreadable and only the examination remains, and the answer must say so rather than manufacture a certainty the bedside does not supply.

4.2 The prescription: the discriminator that comes first

StatPearls sets the drug-class discriminator out cleanly: serotonin syndrome is precipitated only by **agents that impact the serotonergic system**, while neuroleptic malignant syndrome is precipitated by **dopamine receptor antagonists**. The source glosses that second half as the antipsychotics, and taken as a definition the gloss is too narrow in two directions at once. The international consensus work on the syndrome, by Gurrera and colleagues, defines the lesion rather than the drug: a marked reduction in central dopaminergic activity, usually D2 antagonism but also the abrupt withdrawal or reduction of dopaminergic treatment. Read both halves as statements about a mechanism rather than as lists of drugs. Each syndrome is tied to a pharmacological action, not to a therapeutic category and not to a drug name, so the question at the bedside is never whether the patient is on an antidepressant or an antipsychotic by label. It is what the agents on the chart do to the two transmitter systems.

That distinction has consequences the examiner rewards. An agent prescribed for a non-psychiatric indication can still satisfy the serotonergic definition, and an antiemetic or a prokinetic that antagonises dopamine receptors satisfies the other one without ever having been prescribed as an antipsychotic. The withdrawal route leaves nothing on the chart to find at all: there the offending change is a dopaminergic drug that has been stopped or cut, and only the dates will show it. Candidates who scan the drug chart for a diagnostic category miss all of these. Candidates who scan it for a mechanism find them. The receptor pharmacology that underlies these two definitions, taken as a principle rather than as an application, belongs to the Psychopharmacology I: principles and core drug classes notes.

The word *only* in the serotonergic half is load-bearing and cuts in one direction. It closes the causal route in: a patient with no serotonergic exposure at all does not have serotonin toxicity, whatever the picture looks like. It does not work in reverse, because being on a serotonergic agent does not exclude the other syndrome, and a great many patients are on both kinds of drug at once. The class discriminator is therefore a strong exclusion and a weak inclusion, and it has to be used in that direction. That asymmetry is the whole of the difficulty in this item and is taken up below.

4.3 Tempo: the discriminator that costs nothing to obtain

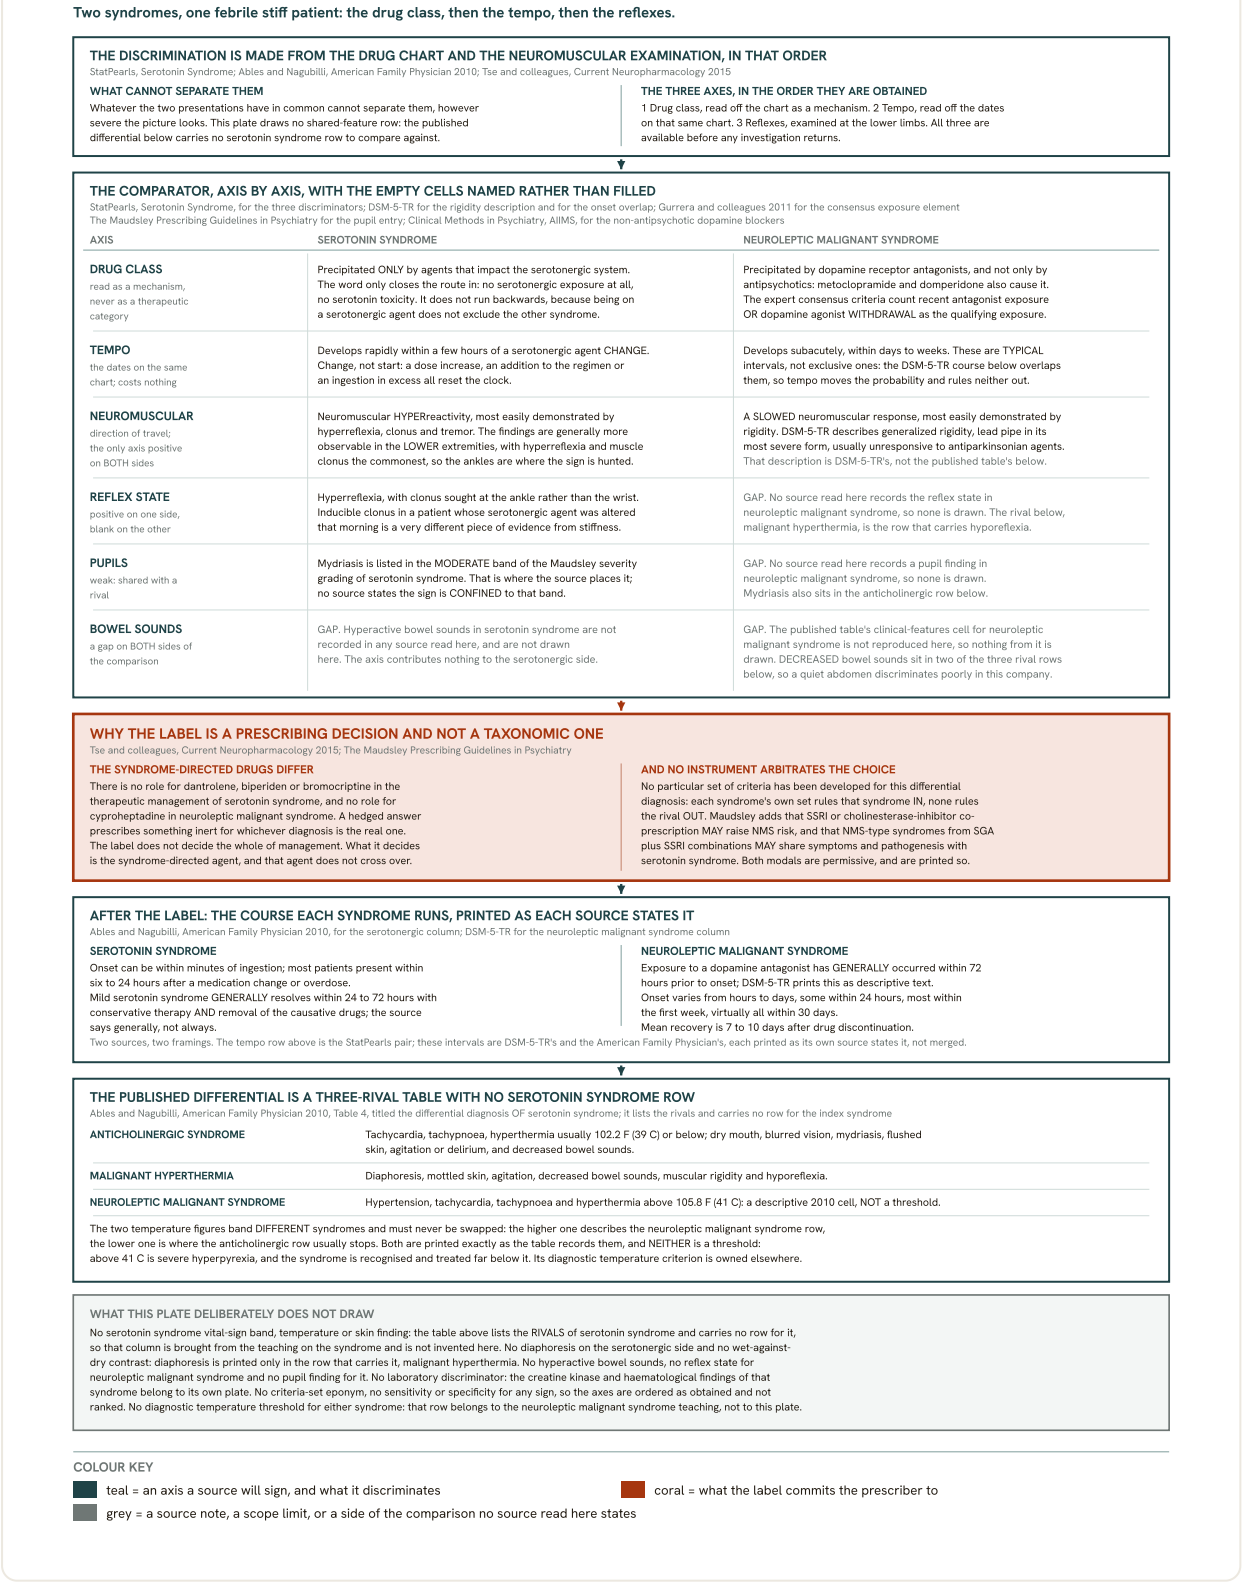


Fig 37.3 Serotonin syndrome against neuroleptic malignant syndrome: the axes that separate them, and the axes that are blank on one side. The three discriminators are obtained in order, drug class then tempo then reflexes, and only the neuromuscular axis returns a positive finding on both sides. Pupils, bowel sounds and the NMS reflex state are drawn as declared gaps rather than filled: the published differential lists the rivals OF serotonin syndrome and carries no row for the syndrome itself, and its two temperature figures band the anticholinergic and NMS rows, never the serotonergic one, and neither is a diagnostic threshold. (StatPearls, Serotonin Syndrome; Ables and Nagubilli, American Family Physician 2010, Table 4; Tse and colleagues, Current

Neuropharmacology 2015; DSM-5-TR; The Maudsley Prescribing Guidelines in Psychiatry; Clinical Methods in Psychiatry, AIIMS.)

The second discriminator is **tempo**, and it is the cheapest one available because the drug chart already carries it. StatPearls times serotonin syndrome to within a few hours of a serotonergic agent change, and neuroleptic malignant syndrome to a subacute development over days to weeks. Those are typical intervals, not exclusive ones. Neuroleptic malignant syndrome can begin far earlier than a subacute description implies, early enough to fall inside the serotonergic window, so the two overlap at their edges. Tempo therefore moves the probability and rules neither syndrome out: clonus hours after a dose change with no antipsychotic exposure is strong evidence, a slow build over a fortnight after an antipsychotic was started is strong evidence the other way, and a stiff, febrile patient on the second day of both drugs is the case tempo cannot settle. The full onset range for neuroleptic malignant syndrome is taught with that syndrome elsewhere in these notes.

Two details make it usable. The first is the word *change*: the trigger is not restricted to a newly started drug, so an increase in dose, an addition to an existing regimen, or an ingestion in excess all reset the clock. The second is that tempo can only be read from a drug chart that carries dates. A history that establishes the patient has been on both an antipsychotic and a serotonergic agent, without establishing when either was last altered, has recorded the exposure and thrown away the discriminator built on it. The accompanying figure sets the tempo axis alongside the other axes of the comparison.

Hours	Days to weeks	Zero	39 C or below
SEROTONIN SYNDROME ONSET, TYPICALLY	NMS ONSET, TYPICALLY	CRITERIA SETS FOR THIS DIFFERENTIAL	ANTICHOLINERGIC ROW, USUALLY

PITFALL

Writing "on olanzapine and escitalopram" in the admission note and nothing more. The entry looks complete and destroys the tempo discriminator, because the reviewer at midnight cannot tell whether one of those agents was changed that morning or has run unchanged for a year. Record the date of the last change to every psychotropic on the chart, not merely the fact of the prescription. The same entry is what a colleague will use to decide whether the picture that develops overnight is the illness or the treatment.

4.4 The neuromuscular examination: opposite directions of travel

The third discriminator is the one that is actually examined, and its strength is that both syndromes give a positive finding rather than one giving a finding and the other giving nothing. StatPearls records the direction of travel as opposite in the two: serotonin syndrome shows **neuromuscular hyperreactivity**, most easily demonstrated by hyperreflexia, clonus and tremor, whereas the examination in neuroleptic malignant syndrome shows a **slowed neuromuscular response**, most easily demonstrated by rigidity.

Stated that way the axis is easy to remember and easy to misapply, because rigidity and hyperreactivity are not mutually exclusive appearances at the bedside. A patient with marked

tremor and increased tone can be described by an inexperienced examiner as stiff. Reflexes and clonus are usually more discriminating than tone alone, which is why the examination has to be performed rather than inferred from the general impression of a rigid, febrile patient. They are not infallible: severe rigidity can obscure them, so an examination that returns no clonus is read alongside the exposure and the tempo rather than banked as a negative result.

Where to look is also specified. The same source records the neuromuscular findings in serotonin syndrome as generally more observable in the lower extremities, with hyperreflexia and muscle **clonus** the most common findings. That is a practical instruction: an upper-limb examination in a restless patient can be equivocal when an ankle examination is not, so the ankles are where the sign is hunted. Inducible clonus at the ankle in a patient whose serotonergic agent was altered that morning is a very different piece of evidence from generalised stiffness in a patient started on an antipsychotic a fortnight ago.

MNEMONIC

Drug class, Tempo, Reflexes

The three axes this comparison actually rests on, in the order they are obtained. The drug class comes from the chart, the tempo from the dates on that chart, and the reflexes from the lower limbs. All three are available before any investigation returns, and none of them requires the patient to be able to give a history.

4.5 The wider differential, and the shape of the published comparison

The two syndromes do not sit alone. The standard published differential of serotonin syndrome, in the American Family Physician review by Ables and Nagubilli, is a three-rival comparison, and its structure teaches something before any of its content does. AAFP Table 4 is titled as the differential diagnosis of serotonin syndrome and its rows are **anticholinergic syndrome**, **malignant hyperthermia** and neuroleptic malignant syndrome. It lists the rivals; it carries no row for the index syndrome itself. A candidate reading it as a two-column contrast will find one column missing, and the serotonin-toxicity side of every axis below has to be brought from the teaching on that syndrome rather than read off the table.

RIVAL IN THE PUBLISHED DIFFERENTIAL	WHAT THAT ROW RECORDS	WHERE THE SEPARATION IS SOUGHT
Anticholinergic syndrome	Tachycardia, tachypnoea and hyperthermia usually **102.2 F (39 C)** or below, with dry mouth, blurred vision, mydriasis, flushed skin, agitation or delirium, and decreased bowel sounds	The mucosa and the skin, which this row records as dry, and a fever band that stops lower than the others
Malignant hyperthermia	Diaphoresis, mottled skin, agitation, decreased bowel sounds, muscular rigidity and hyporeflexia	The exposure history, and reflexes that are depressed rather than brisk beside rigidity
Neuroleptic malignant syndrome	Hypertension, tachycardia, tachypnoea and hyperthermia above **105.8 F (41 C)**, which is how this 2010 table describes the row and is not a diagnostic threshold	The prescribed class and the interval since it changed, before any physical sign. The temperature is not the separator here, and the syndrome is diagnosed well below this figure

Three points earn marks off that table. Mydriasis appears in the anticholinergic row, so a dilated pupil is not by itself a serotonergic sign. Decreased bowel sounds appear in two of the three rows, so a quiet abdomen discriminates poorly in this company. And the two temperature figures band different syndromes and must not be swapped: the higher one describes the neuroleptic malignant syndrome row, the lower one is where the anticholinergic row usually stops. Neither figure is a diagnostic threshold, and the higher one is dangerous read as one. Above 41 C is severe hyperpyrexia; neuroleptic malignant syndrome is recognised and treated far below that, so a candidate who waits for the table's number before naming the syndrome has misread a descriptive cell as a criterion. The diagnostic temperature criterion is taught with that syndrome elsewhere in these notes and is deliberately not printed here. Anticholinergic toxicity, lithium toxicity and malignant catatonia are taught as entities in their own right elsewhere in these notes; what belongs here is only their role as rivals in this one comparison.

4.6 The criteria question, and what no instrument will do for you

Candidates reach for a criteria set at exactly this point, and the reach is the trap. Each syndrome has a named criteria set of its own, taught with that syndrome elsewhere in these notes, and each was constructed to identify its own syndrome in a patient suspected of having it. An instrument built to rule one condition in is not, by construction, an instrument that rules a rival out. The review in Current Neuropharmacology states the position for this pair directly: no particular set of criteria has been developed for this differential diagnosis.

TRAP **Quoting a criteria set as though satisfying it excluded the other syndrome.** A patient exposed to both drug classes can satisfy a serotonin-toxicity criteria set and still have the other diagnosis, because those criteria were validated to identify serotonin toxicity and not to exclude the rival syndrome, and a mixed exposure can produce findings that satisfy them anyway. The examiner is testing whether the candidate knows that the published instruments are diagnostic aids for one syndrome each, and that none has been validated against the specific problem of choosing between the two.

TRAP **Treating the two syndromes as points on one hyperthermic spectrum and hedging the management.** Candidates who cannot decide sometimes write that they would cover both. The published position is the opposite of a spectrum: the syndrome-directed agents used in one have no role in the other, so a hedged answer adds a drug that is inert for whichever diagnosis is real, and it tells the examiner the candidate never made the discrimination. Stopping the suspect agent and supporting the patient is not the hedge, because that is shared ground. Prescribing both syndrome-directed drugs is.

4.7 The patient who is on both drug classes

This is the case the item exists for, and it is the case the discriminators handle worst. The Maudsley Prescribing Guidelines state that co-prescription of SSRIs or cholinesterase inhibitors with antipsychotic medication may increase the risk of neuroleptic malignant syndrome, and that NMS-type syndromes induced by combinations of second-generation antipsychotic and SSRI medications may share their symptoms and pathogenesis with the serotonin syndrome. The modal in both statements is permissive and must be reproduced as such: this is a stated possibility about risk and about shared mechanism, not an established equivalence between the two syndromes.

Hold that beside the criteria position and the picture is coherent rather than merely unsatisfying. The combination that is hardest to classify is also the combination in which the two syndromes are suggested to overlap in mechanism, and it is precisely the combination for which nobody has built a discriminating instrument. The honest teaching is that in the co-prescribed patient the discrimination rests on the drug history and the physical examination, and that a confident label offered without either is a guess wearing a diagnosis.

IN INDIA

Antipsychotic and SSRI co-prescription is routine practice here, so the ambiguous case is the ordinary case rather than the examination curiosity. Two decisions follow. Date every change to both drugs in the notes, because co-prescription is what weakens the drug-class axis and tempo is one of the two discriminators left standing when a patient arrives already on both. The other is the neuromuscular examination, and both have to be recorded. And examine the ankles for reflexes and clonus in every febrile, stiff patient on a combined regimen, because that examination is free, is available in any setting, and is the one axis that gives a positive finding on both sides of the comparison. Do not wait for a criteria set to arbitrate, because none has been developed for this differential.

4.8 Writing the comparison, and what the label commits you to

A comparative-differentiation answer is marked on structure. Four elements carry it, and they are obtained in this order.

- The pharmacological class of every agent on the chart, read as a mechanism rather than as an indication, because that is the discriminator the source puts first.
- The date of the last change to each of those agents, which is what makes the tempo axis readable at all.
- The neuromuscular examination, directed at the lower limbs, looking for the direction of travel rather than for stiffness alone.
- The named diagnosis, followed by the statement that the syndrome-directed drug for one has no role in the other, which is what makes the naming consequential.

On management the boundary holds. The locus of care and the escalation belong to the Perinatal/women's mental health and emergency psychiatry notes. What the discrimination decides is the focus and the mode: the offending agent to be stopped is identified by the class definition above, and the specific agent added afterwards belongs to whichever syndrome has been named, since the review is explicit that neither syndrome's named agents have a role in the other. Each syndrome's own management, in full, is taught with that syndrome elsewhere in these notes and is not repeated here.

GOOD TO REMEMBER

When the diagnosis is genuinely undecided in a co-prescribed patient, what is not in dispute is withdrawal of the suspect agents and the supportive care around that withdrawal, because those steps are common to both syndromes while the syndrome-directed drug that follows is specific to one of them. Deciding which agent is suspect returns you to the class definitions, which is why those two definitions are worth knowing verbatim rather than approximately.

The prescription and the reflexes decide this comparison, the tempo shifts it without settling it, and the decision commits the patient to one of two syndrome-directed drugs, neither of which has a role in the other syndrome.

5 · Other psychotropic emergencies (anticholinergic toxicity, lithium toxicity, malignant catatonia)

Two of these three states are made by the prescription; the third is made by the illness and then complicated by it. Anticholinergic toxicity and lithium toxicity are drug toxicities. Malignant catatonia is not: it is a severe catatonic state usually complicating a psychiatric or medical disorder, and the prescription enters as a competing explanation rather than a cause. The three are grouped because in each of them the drug history does work that no bedside sign can do: it names or excludes a mechanism, it fixes what a measured concentration is allowed to mean, and it decides what happens to the prescription afterwards. The emergency-department presentation, the resuscitation and the disposition belong to the Perinatal/women's mental health

and emergency psychiatry notes, and the receptor logic behind each drug class belongs to the Psychopharmacology I: principles and core drug classes notes. What is examined here is the pharmacology in between: which agents and combinations produce the state, what a serum value can and cannot settle, which prescriptions carry a stable patient across the line, and what the standard sources actually say about treating each one.

5.1 The anticholinergic toxidrome, and the mnemonic that does not quite fit

Muscarinic blockade produces a picture stereotyped enough to be taught as a rhyme. StatPearls sets out the **anticholinergic toxidrome** in six phrases, and then lists the signs they stand for.

MNEMONIC

red as a beet, dry as a bone, blind as a bat, mad as a hatter, hot as a hare, full as a flask

The signs the same source pairs with them are flushing, anhydrosis, dry mucous membranes, mydriasis, altered mental status, fever and urinary retention, with decreased bowel sounds a further common finding. Count the two lists against each other before reciting them: there are seven signs against six phrases, so the one-to-one correspondence the source implies does not fully resolve, and the rhyme is a memory aid rather than a mapping.

Central involvement is what turns this from an unpleasant side-effect burden into an emergency. The same account gives delirium, hallucinations, agitation, restlessness and confusion, together with the two features it treats as characteristic of this particular delirium, **staccato speech** and picking at clothing and bedding, and notes that seizures and jerking movements are possible. The peripheral half of the syndrome tells you which receptor is blocked; the central half tells you the patient is in trouble.

The pharmacological point that organises all of it is that the offending agents are seldom labelled anticholinergics. Tricyclic antidepressants, low-potency first-generation antipsychotics, clozapine, the antiparkinsonian anticholinergics prescribed for extrapyramidal side effects, sedating antihistamines and a long tail of preparations bought without prescription all contribute, and the burden is additive across them. Toxicity may follow a single agent taken in overdose, and it may equally be an arithmetic problem across a whole prescription with no one drug on it at fault, which is why the drug history has to be taken by class of action and not by indication.

5.2 Benzodiazepine first, antidote second

The treatment sequence is examined more often than the toxidrome and it is counter-intuitive. StatPearls makes intravenous benzodiazepines the first-line therapy for agitation, potentially in large doses, in anticholinergic poisoning, and does not put the antidote first.

RULE **Treat the agitation before you treat the receptor.** The sequence is StatPearls' own, and what recommends it is that sedation stakes less on the diagnosis being right than a specific antidote does. A benzodiazepine is not thereby harmless: in the doses this picture needs it can deepen the confusion and compromise breathing. Whether the antidote follows turns on diagnostic confidence, the electrocardiogram and the contraindications, a toxicology decision this account does not make. A stem describing an agitated, flushed, dry, urinary-retaining patient is testing whether the candidate sedates first and reserves the antidote for the case that earns it.

Physostigmine is the specific antidote, a cholinesterase inhibitor whose effect is felt centrally as well as peripherally, and the source notes that it is usually only given in the case of both peripheral and central signs and symptoms of anticholinergic poisoning. That is a description of practice and not an eligibility rule, though it does encode a differential test as much as a treatment threshold. A purely peripheral picture seldom calls for a centrally acting drug, and a central picture with no peripheral signs at all is a reason to reopen the differential rather than to reach for the antidote.

The figures are worth carrying, with the caveat attached to them. StatPearls gives the recommended physostigmine dose in anticholinergic poisoning as **0.5 to 2 mg IV for adults and 0.02 mg/kg IV for paediatric patients, with a maximum paediatric dose of 0.5 mg**. Read that as a stated dose with a paediatric ceiling rather than as a full regimen: no dosing frequency and no adult maximum are stated with it, no titration schedule is carried, and no separate initiation dose sits below it, because the source prints one recommended dose rather than the first step of a schedule. It states no administration rate either, and rate is not a detail: Dawson and Buckley's review of anticholinergic delirium describes physostigmine as titrated intravenously under monitoring, because speed of administration is where the cholinergic and cardiac risk sits. No rate is carried here, and the figure above is not licence to give the dose in one push. That is an internationally recommended dose from a source that carries no Indian regulatory or formulary statement with it, so whether a given unit here can obtain the drug at all is a local question this account does not answer.

5.3 Lithium toxicity: the number, and the band it sits above

Lithium toxicity is the psychotropic emergency tied most closely to a laboratory value, but the value supports the diagnosis rather than making it: exposure pattern, renal function, clinical state and the direction of serial concentrations settle both the toxicity and what is done about it. Meyer and Stahl's Lithium Handbook records that in contemporary usage toxicity is often defined as a serum concentration of **1.5 mEq/L or above**. The Maudsley Prescribing Guidelines describe the same territory in bands: above 1.5 mmol/L toxic effects reliably occur and usually consist of gastrointestinal effects, increasing anorexia, nausea and diarrhoea, and central effects, muscle weakness, drowsiness, confusion, ataxia, coarse tremor and muscle twitching, while above 2 mmol/L increased disorientation and seizures usually occur, which can progress to coma and ultimately death. The hedge is the source's own: both bands are stated as what usually happens. The same section then limits its own figures, noting that these plasma concentrations are only a guide, and individuals vary in their susceptibility to symptoms of toxicity.

What makes the threshold examinable is how close it sits to the top of the treatment band. The Indian Psychiatric Society bipolar guideline states in its narrative that target serum concentrations for acute-phase efficacy fall between 0.6 and 1.2 meq/l. The distance between the top of that band and the concentration at which toxicity is defined is small, and that narrowness, not the absolute value of either figure, is the argument for measuring the drug at all. The published target ranges for the individual agents are taught with therapeutic drug monitoring; what belongs here is only where the band stops.

TRAP **Quoting one guideline's lithium target as though the guideline held only one.** The Indian Psychiatric Society bipolar guideline disagrees with itself. Its narrative adds that 0.4 to 0.8 meq/l may suffice for older adults, which is a permissive statement rather than a target; its own episode-specific dosing tables give a different older-adult figure again, and no erratum reconciles the two loci. A candidate who quotes either one alone is wrong against the other part of the same document, so name the disagreement instead of resolving it.

5.4 Three exposure patterns, and why the same value is not the same event

A serum concentration measures the compartment that is easy to sample, not the one that matters. The Lithium Handbook classifies toxicity by exposure pattern into acute, **acute on chronic** and chronic, and records that the latter two are the more serious in nature and account for the great majority of overdose or toxicity presentations. It gives the mechanism in the same chapter: patients chronically treated with lithium already carry brain concentrations high enough that an acute ingestion, or prolonged exposure to toxic levels, typically 2.0 mEq/L or above, is associated with poorer outcomes than comparable exposures in a lithium-naive patient.

RULE **Ask how long, not only how much.** The identical concentration in someone who swallowed a bottle yesterday and in someone who has taken lithium for a decade describes two different brain exposures, and the sources rank the chronic and acute-on-chronic patterns as the more serious ones. Duration of exposure is part of reading the value, not background to it.

Against that, the routine case has become uncommon. The Lithium Handbook records that modern monitoring schemes and serum concentration ranges have significantly lowered rates of lithium toxicity, and gives an incidence figure reported in longitudinal studies, with a Swedish cohort named alongside it. No comparable Indian incidence figure is available, so the rarity should be read as a property of well-monitored practice rather than as a property of the drug.

5.5 When the concentration is not the syndrome

The most dangerous lithium result is a normal one. The Maudsley guidelines record that neurotoxicity at normal plasma concentrations has also been described, because brain lithium concentrations may not be reflected in the plasma. That one sentence dismantles the reflex that a value inside the range excludes toxicity, and it is why this diagnosis is clinical with laboratory support rather than the reverse.

The severe end of that spectrum has a name. **SILENT**, the syndrome of irreversible lithium-effectuated neurotoxicity, denotes neurological sequelae that persist after the drug has been withdrawn, the original description setting the boundary at sequelae still present at least two

months after lithium is stopped. Persistent cerebellar dysfunction is the commonest pattern, not the whole of the definition. A review of published cases found that among the subset of 51 cases with persistent cerebellar symptoms whose peak concentrations were under 2.5 mEq/l, fever or infection was a strong predictor, with an odds ratio of 13.9 (95% CI 3.21 to 60.05), while in the 34 cases whose peak concentrations reached or exceeded that figure fever was not a significant predictor of cerebellar sequelae. Read the scope carefully. This is a review of published cases, and 2.5 mEq/l is that review's cut, not a validated line above which a patient is in danger and below which the numbers reassure. Within that selected material, fever or infection carried the association with persistent cerebellar sequelae in the subgroup under the cut, and not in the subgroup above it.

PITFALL

Sending the ataxic patient home because the level came back in range. Brain and plasma are not the same compartment and the source describes neurotoxicity at normal plasma concentrations, so new coarse tremor, ataxia or confusion in a lithium-treated patient is a neurological problem that needs an explanation whatever the assay says. Where a febrile illness sits behind it, the fever is not a side issue to the lithium question; in the reviewed cases whose peak concentrations stayed under the review's own cut, it was the variable that carried the association with persistent cerebellar sequelae.

5.6 The prescriptions, and the sample, that anticipate it

Lithium offers no receptor profile from which to predict its adverse effects. What predicts them is exposure: how much drug, for how long, and what is happening to its clearance. The interactions that matter are therefore the ones that alter renal handling of sodium and water, and they come from general medicine rather than from psychiatry. The Maudsley guidelines report that in the elderly, angiotensin-converting enzyme inhibitors increase seven-fold the risk of hospitalisation due to lithium toxicity. For thiazide diuretics the timing is predictable and the size of the effect is not: concentrations usually rise within 10 days of the diuretic being prescribed, and the magnitude of that rise varies widely between patients, which means no safe assumption can be made about how far a given patient will move.

The Indian Psychiatric Society bipolar guideline converts the same pharmacology into a monitoring instruction, naming the patients who need **more frequent monitoring**.

- Older adults.
- Those receiving diuretics.
- Those receiving non-steroidal anti-inflammatory drugs or angiotensin converting enzyme inhibitors.
- Those with renal impairment.

None of it is interpretable without the sampling rule attached to it. The same guideline requires that the lithium sample be drawn **11 to 13 hours after the last dose**, and before the next dose where the drug is given twice or three times daily. That window belongs to routine monitoring: a

concentration drawn outside it cannot be read against the monitoring bands and will not carry a dose change. It is not a reason to withhold a level from a patient who looks toxic. In suspected poisoning the sample is taken on presentation and repeated, and the interpretation rests on that trend, the time since the last dose and the clinical state.

>90%	48%	25 to 400%	0.01
CHRONIC OR ACUTE ON CHRONIC	REVIEWED SILENT CASES WITH FEVER OR INFECTION	RISE AFTER A THIAZIDE IS ADDED	TOXICITY EPISODES PER PATIENT-YEAR

IN INDIA

Two outpatient decisions follow from that footnote and that sampling rule. First, ask about analgesics by name at every review rather than asking whether the patient takes any other medicines: an anti-inflammatory bought across a counter is not experienced as a prescription and will not be volunteered, yet the guideline names that class alongside diuretics and angiotensin converting enzyme inhibitors as a reason to measure more often. Second, book the blood draw to the guideline's own post-dose window rather than to whatever hour the patient reaches the clinic, because a conveniently timed sample cannot support a dose change in either direction. The unwell patient is the exception noted above: there the blood is drawn on arrival and repeated, not deferred to the next morning's window.

5.7 Malignant catatonia, and the syndrome it may not be separable from

Malignant catatonia is the third emergency and the hardest to place, because the authorities that describe it say it may not be separable from something else. The Maudsley Prescribing Guidelines state that it may not be distinguishable from neuroleptic malignant syndrome, either clinically or by laboratory testing, and that this has prompted the view that the latter may be a variant form of the former. DSM-5-TR says the same thing from the opposite direction, in its differential diagnosis of neuroleptic malignant syndrome: malignant catatonia may be indistinguishable from it, and some investigators consider neuroleptic malignant syndrome to be a drug-induced form of malignant catatonia. Two standard sources converging on the same non-separation is not a gap in the literature to be filled in an answer; it is the finding.

That leaves the drug history as the most useful discriminator, and the sources are precise about how far it reaches. In the absence of any prior or recent administration of a **dopamine antagonist**, the Maudsley position is that neuroleptic malignant syndrome can probably be ruled out. The hedge is the source's word and is not to be upgraded, and the exclusion is about antagonist-associated disease: the fall in central dopaminergic activity that produces the syndrome also follows abrupt withdrawal of dopaminergic treatment, which no drug chart records as an exposure. The drug history shifts the probability; it does not settle the diagnosis. What the source does not state is the converse.

TRAP **Reading the exclusion rule backwards.** One inference is licensed, and only along the antagonist route: no dopamine antagonist, probably not neuroleptic malignant syndrome. The converse, that recent exposure makes the diagnosis, is not licensed, and the converse is what candidates write. A catatonic patient who has also received an antipsychotic has two candidate states and one shared laboratory picture, and naming the exposure does not choose between them.

The exposure list is also wider than the antipsychotics. The Maudsley guidelines record that neuroleptic malignant syndrome is also sometimes seen with antidepressants, valproate, phenytoin and lithium. That matters twice in this topic: it widens the prescription review that has to be done before the syndrome is excluded on drug-history grounds, and it puts lithium on both sides of the same page, as an agent that produces its own concentration-defined emergency and as an agent occasionally implicated in this one.

5.8 Pharmacological management, and the marker the two share

Malignant catatonia is not defined here by a criteria set, and that absence is part of the teaching. The standard guidelines and textbooks describe it as a clinical state and say that it may not be distinguishable from neuroleptic malignant syndrome, either clinically or by laboratory testing; none of them supplies a threshold that makes catatonia malignant. An answer that produces criteria for it has invented them. The rating instrument for catatonia and the challenge test that establishes the diagnosis belong with the catatonia teaching rather than with this emergency, and are not redrawn here.

Treatment, in the sources read for this topic, is stated for the refractory case only. The Indian Psychiatric Society schizophrenia guideline places dantrolene and dopamine agonists like bromocriptine in refractory malignant catatonia. Read the qualifier rather than the drug names: refractory. These agents are named for the case that has not responded, and the guideline does not present them as the opening pharmacological move. The emergency act itself, including escalation and the disposition decision, belongs to the Perinatal/women's mental health and emergency psychiatry notes.

EMERGENCY	WHAT PRODUCES IT	WHAT A MEASUREMENT CONTRIBUTES	THE COMMONEST WRONG INFERENCE
Anticholinergic toxicity	A single agent in overdose, or additive antimuscarinic burden across prescribed and over-the-counter agents	No defining concentration; the peripheral and central halves of the syndrome are what gate the antidote	That the specific antidote is the opening move
Lithium toxicity	Acute ingestion, or accumulation from impaired renal clearance, volume or sodium depletion, a dosing error or an interacting drug	Support, not a verdict: read against the bands only on a correctly timed sample, and in suspected toxicity it is the serial trend with the exposure pattern that counts	That a value inside the treatment range excludes neurotoxicity
Malignant catatonia	Usually a severe catatonic state arising in psychiatric or medical illness, not a drug effect; antipsychotic exposure raises the competing possibility of neuroleptic malignant syndrome and may worsen the catatonia	Nothing discriminating; the laboratory picture is shared with neuroleptic malignant syndrome	That recent dopamine antagonist exposure settles the diagnosis

GOOD TO REMEMBER

Kaplan and Sadock's Comprehensive Textbook of Psychiatry supplies the one laboratory value the two overlapping states share: serum **creatine kinase** is especially elevated in malignant catatonia and in neuroleptic malignant syndrome and can increase the risk of renal failure, and should be tracked to ensure a downtrend. The instruction is serial, not single. One raised value describes a moment; a falling series is one objective marker that the decision already taken is working, read alongside temperature, autonomic stability, rigidity and renal function rather than in place of them, and the cheapest of them in any unit.

In all three emergencies the prescription is where the reading starts and not where the diagnosis ends: read the exposure and its duration first, then the number, never the number by itself, and remember that the third of them can arrive in a patient whose drug chart had nothing to do with it.

Treatment resistance and augmentation

6 · Treatment-resistant depression: definition, staging and assessment

Resistance is a judgment on the treatment record before it is a statement about the patient. Almost every element of the definition of treatment-resistant depression is a convention that

somebody wrote down: how many failed trials, how large a fall in symptoms counts as a response, how long a trial must run before it may be called failed, and whether the failures have to be pharmacological at all. Different bodies settled those questions differently, and the better ones say so in the same paragraph as their own definition. Marks here go to an answer that states a definition and names whose it is, shows what has to be excluded before the label is applied, knows the response arithmetic sitting underneath it, and can say why the staging models built to grade resistance never entered routine practice. It is examined as a definition question and it is answered badly as one, because the definition is the least interesting part of the topic.

6.1 The definition, and the bodies that issue it

The most widely used definition of **treatment-resistant depression** is inadequate response to **a minimum of two antidepressants** given at **adequate dose and duration** with adherence confirmed, and it carries regulatory weight because both the United States Food and Drug Administration and the European Medicines Agency have adopted it. The Indian CPG 2025 sets the same bar, at least two antidepressants at adequate dose and duration, then qualifies it in the same breath: definitions vary, and they commonly overlook psychological and neurostimulation treatments and partial responses. CANMAT 2023 is blunter, recording that a universal consensus definition is lacking and that the most frequently used one is failure to respond to at least two antidepressant trials at a therapeutic dose and adequate duration. One caveat attaches to that third attribution and not to the bar itself: the CANMAT sentence quoted here could not be located in the copy of the 2023 update used for these notes, so treat the CANMAT wording as unverified and rest the two-antidepressant threshold on the regulatory and Indian statements, which are. Three bodies, one arithmetic core, and an open admission from two of them that the core is an agreement rather than a finding.

The operational detail is where they separate. The accompanying table puts the same four questions to each definition, and the blanks are as informative as the entries: a definition that fixes a trial count without fixing a non-response threshold has left the harder half of the work undone.

DEFINITION	FAILED TRIALS REQUIRED	THRESHOLD FOR NON-RESPONSE	MINIMUM TRIAL LENGTH STATED
Food and Drug Administration and European Medicines Agency, as the most used definition	A minimum of two antidepressants, with adequacy of the trial and adherence both required	Not stated in the definition	Not stated in the definition
Indian CPG 2025	At least two antidepressants at adequate dose and duration	Not attached to the definition	Not attached to the definition
CANMAT 2023	At least two trials at a therapeutic dose and adequate duration	Not stated with the definition	Not stated with the definition
Canadian expert consensus, by modified Delphi	Fixed by the consensus; figure not printed here	Fixed by the consensus; figure not printed here	Fixed by the consensus; interval not printed here
European Group for the Study of Resistant Depression	Consecutive adequate trials, characterised by named clinical predictors; the count and the predictors are not printed here	Below 50 percent symptom reduction	Not stated in the account used here

Only the Canadian consensus fixes all four, and its own report states the level of agreement each of those items reached at GRADE-high evidence. None of the four values, and none of the agreement figures, is printed in this account. The Delphi report itself could not be consulted here, and every one of those numbers is one a reader would apply to a patient's trial record, so the table says which columns the consensus filled and leaves the values to the paper. What the comparison still shows is the shape of the problem: a formal vote was needed to fix elements the other definitions leave open, which is a fair measure of how soft the ground is. The European group adds something the others leave out: it characterises its trials by the clinical features that travel with resistance rather than by the trials alone, and it operationalises response with validated rating scales instead of clinical impression. How many consecutive trials it requires, and which clinical features it names, both come from a table this account could not check against the original report either, so the requirement is described as a class and its contents are not listed.

RULE **Treatment resistance is an operational judgment, not a property of the illness alone.** It combines documented adequate treatment exposure, meaning how many agents, at what dose, for how long, and whether they were swallowed, with inadequate clinical response in a depressive illness. The exposure half is entirely a matter of record, which is why the first move on meeting an apparent non-responder is to audit the treatment record rather than escalate the regimen, and why the guidelines quoted here prescribe a reappraisal before the next prescription.

6.2 What makes a trial adequate

The trial count is arithmetic; the word carrying the weight is adequate, and it has three limbs. Duration is the limb most often assigned a number, but the published thresholds are not uniform. Across the literature the minimum antidepressant trial duration cited for adequacy is typically 4 weeks, with a reported range running from 4 to 12 weeks. The Indian material sits at the shorter end and is stable across two editions of guidance: clinical improvement with an antidepressant typically appears after **4 to 6 weeks** of adequate treatment, and the IPS CPG 2017 depression guideline states that adequate treatment for at least 4 to 6 weeks is necessary before concluding that a patient is not responsive to a particular medication. The Canadian expert consensus voted on a minimum of its own, and for the reason given above that figure is left to its paper. The Indian interval, stable across two editions of guidance and matching the shorter end of the published range, is the one to carry into the hall.

Dose adequacy resists a single number because it is agent-specific: a therapeutic dose of one antidepressant is a fraction of the daily milligram figure of another, which is why the definitions ask for a therapeutic dose of that agent and not a universal one. Which agent is chosen within a class, and the pharmacokinetic reasoning that fixes what a therapeutic dose is, belong to the Psychopharmacology I: principles and core drug classes notes. The third limb, adherence, is written into the most used definition itself, which is why it has to be confirmed rather than assumed before a trial is counted as failed.

6.3 How much improvement counts as a response

A trial can only fail against a threshold, and the conventional set is straightforward: a **partial response** is at least 25 percent but less than 50 percent improvement and **non-response** is **less than 25 percent improvement**. The 50 percent mark is where a response begins, so the three categories partition the range without overlapping. The resistance literature works from the upper of those two boundaries, not the lower. The European Group for the Study of Resistant Depression identifies resistance as below 50 percent symptom reduction, and the Canadian consensus voted a threshold of its own that this account does not print. A patient who is a partial responder on the conventional set can therefore be counted a non-responder by the resistance definitions, and the two words are not interchangeable in a written answer.

TRAP **Reading a response-monitoring rule as an adequacy criterion.** The IPS CPG 2025 depression guideline directs that where less than 50 percent improvement is observed even after 6 to 8 weeks of treatment at optimal dose with good adherence, treatment is modified by either switching or augmenting. That is a decision rule for a trial already under way, and it belongs to the guideline's material on how long an antidepressant is continued, not to its account of resistance. The section defining resistance asks only for an adequate dose and duration and attaches no number to either. A candidate who converts the monitoring interval into the adequacy threshold has printed a figure the guideline never issued as one.

6.4 Pseudo-resistance, and the work that precedes the label

The first question to put to an apparent non-responder is whether the resistance is real. A substantial proportion of apparent treatment resistance is **pseudo-resistance**, arising from

inadequate trials or non-adherence rather than true refractoriness. The distinction changes the action completely: one situation calls for finishing a trial that was never finished, the other for moving past it. Indian guidance puts the audit first, stating that the first step in caring for a patient who has not responded to medication is a thorough review and reappraisal of the psychosocial and biological information base, aimed at re-verifying the diagnosis and identifying neglected contributing factors. That places a review, not an escalation, at the head of the algorithm.

MNEMONIC

Diagnosis, Adherence, Trials, Causes, Stressors

Five checks stand between a non-responder and the label. The first three interrogate what was prescribed and whether it was taken; the last two ask what else is holding the episode open. Work them in that order, because a diagnostic revision makes the trial history irrelevant, whereas a stressor identified first tempts the clinician to stop looking.

IN INDIA

Before confirming resistance, the Indian 2025 guideline requires five assessment steps: reassess the diagnosis, with bipolar disorder and comorbidity specifically in mind; check treatment adherence; review the adequacy of past trials for dose, duration and switching strategy; identify medical causes, naming hypothyroidism and vitamin deficiencies; and evaluate psychosocial stressors. Those two are given as examples of the medical causes to look for, not as a panel to run on every non-responder: investigate thyroid dysfunction or nutritional deficiency where the clinical assessment and the local protocol indicate it. The step sits before the label in the guideline's own order, so the question is settled while resistance is being considered rather than after it has been recorded.

PITFALL

Counting a trial the patient abandoned in the second week as one of the two failures. The trial count is the easiest part of the definition to satisfy on paper and the easiest to satisfy wrongly. An agent stopped early for side effects tested tolerability, not efficacy; an agent taken intermittently tested neither. For every past drug, establish the dose reached, the time spent at it, and how much of it was actually taken, before that drug is allowed into the count.

6.5 How common resistance is, and what each figure counts

Two prevalence figures are routinely quoted against one another and they count different things. Per the Indian 2017 guideline, initial treatment with an antidepressant fails to achieve a satisfactory response in approximately 20 to 30 percent of patients with depressive disorder, which is the failure rate of the first trial. Measured against the two-trial definition, **at least 30 percent** of people with depression meet criteria for resistance. The two numbers look alike and answer different questions, so setting them side by side as an Indian figure against a Western one is an error about denominators rather than about arithmetic. No Indian prevalence estimate for

resistance itself was available for this account; the defensible answer in a viva is to give the Western figure and label it as Western.

Attrition across successive treatment steps is the more useful teaching number, because it quantifies what each further attempt is worth.

36.8%	30.6%	13.7%	13.0%
FIRST STEP	SECOND STEP	THIRD STEP	FOURTH STEP

In STAR*D, QIDS-SR remission rates fell across the four successive acute treatment steps as shown, with an overall cumulative remission rate reported at 67 percent. The collapse between the second step and the third is the single most quotable finding in this literature, because it marks where the yield of simply trying another agent falls away, and it is the empirical argument for treating two failures as a decision point rather than an arbitrary one. The cumulative figure is a different matter.

TRAP **Quoting the cumulative remission figure as a settled result.** A reanalysis of the original STAR*D data found the overall remission rate to be 35 percent, as opposed to the claimed 67 percent, and the standard prescribing texts now print the reanalysis alongside the original. The step-wise rates are not what is contested; the cumulative arithmetic is. Quote the step-wise collapse freely, and never quote the cumulative figure without its reanalysis beside it.

6.6 The staging models, and what each of them grades

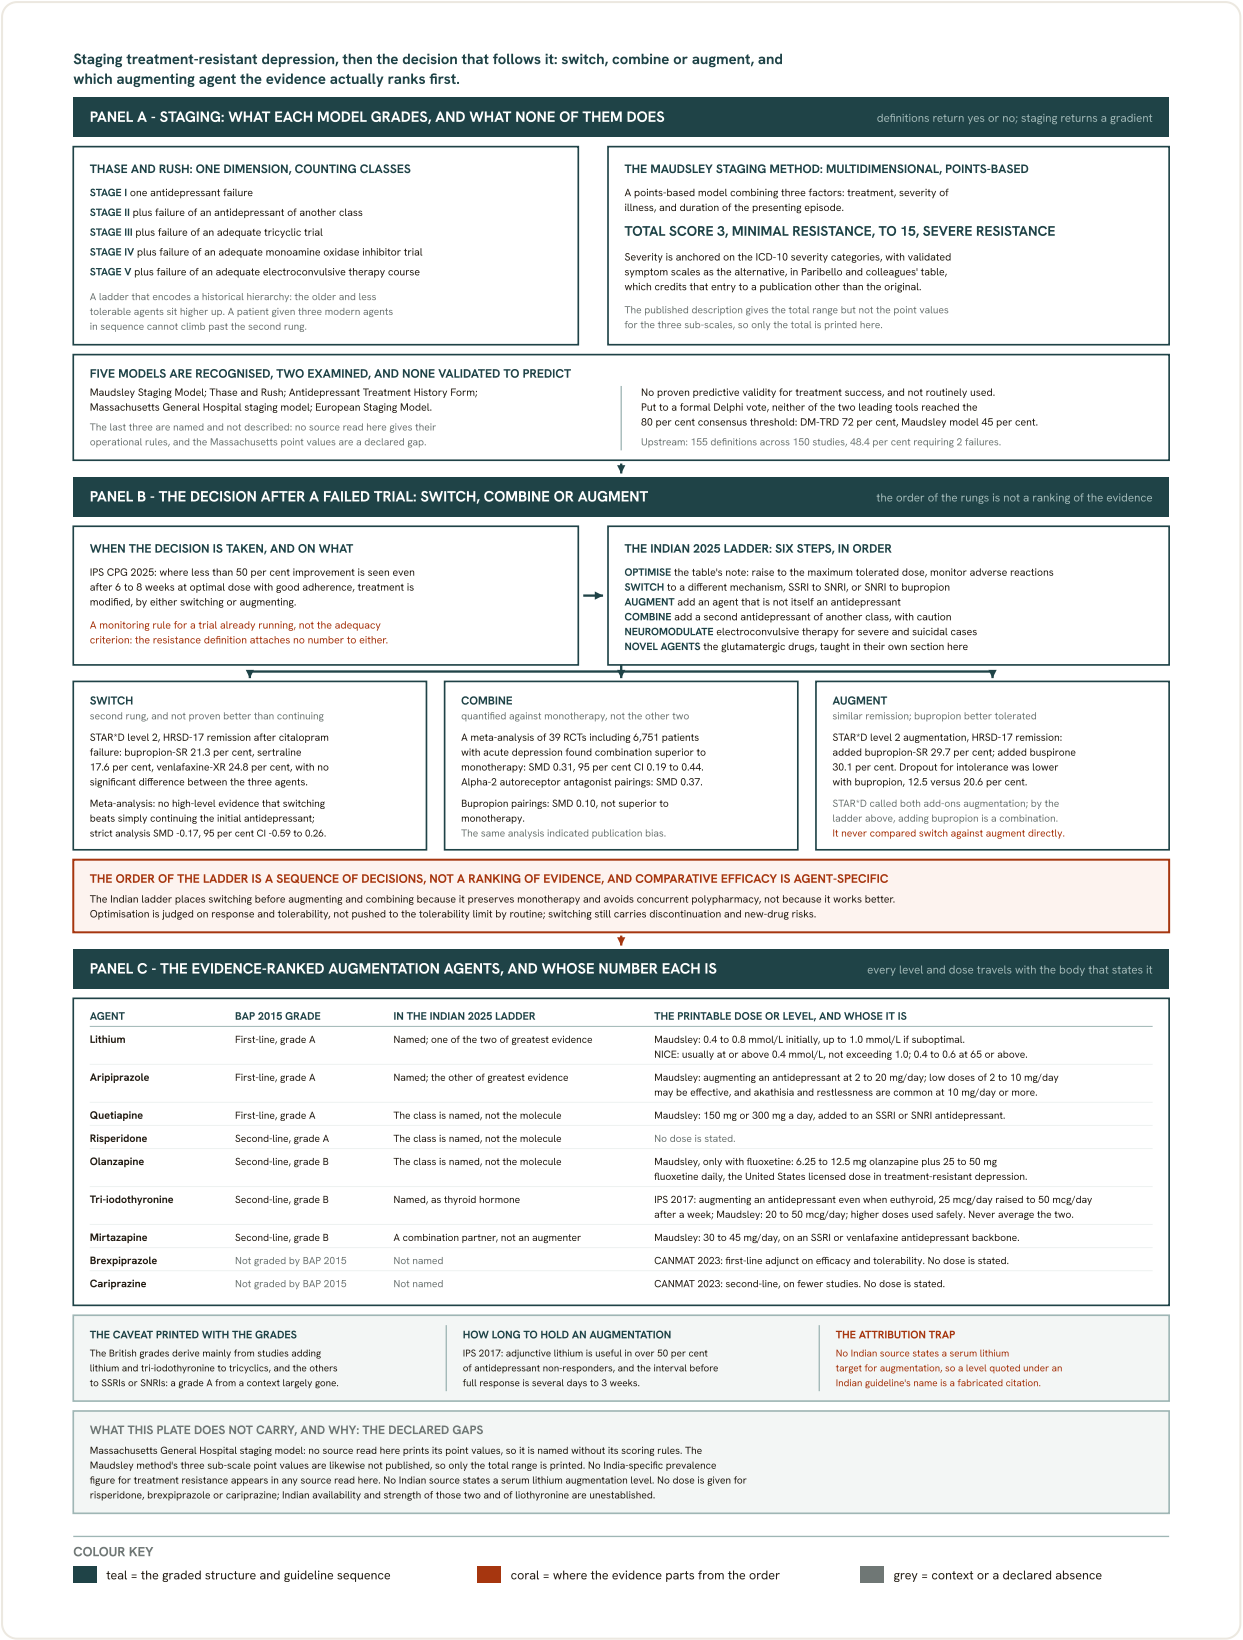


Fig 37.4 Staging treatment-resistant depression, then switching, combining or augmenting. Two of five recognised models are examined: Thase and Rush counts failed classes across five stages, the Maudsley method scores treatment, severity and duration to 3 to 15. Neither is validated to predict, and at a Delphi vote neither cleared 80 per cent. The Indian 2025 ladder runs optimise, switch, augment, combine, neuromodulate, novel agents, a decision sequence, not a ranking of evidence: switching is early because it preserves monotherapy, optimisation is the table's wording qualified in the rule box, and comparative efficacy is agent-specific. Every dose and level is attributed to the body stating it; the Indian guideline names the agents without a figure. (IPS CPG

2025 and 2017; Maudsley Prescribing Guidelines 15th edition; NICE NG222; BAP 2015; CANMAT 2023; STAR*D; Rybak and colleagues 2021; Paribello and colleagues 2026.)

Where a definition returns a yes or a no, **staging models** return a gradient. Five are recognised:

- the Maudsley Staging Model;
- the Thase and Rush model;
- the Antidepressant Treatment History Form;
- the Massachusetts General Hospital staging model;
- the European Staging Model.

Two of them are worth knowing in detail, because they grade different things and a question asking you to compare staging systems is asking for exactly that contrast. The **Thase and Rush model** is one-dimensional and counts classes: Stage I is a single antidepressant failure; Stage II adds failure of an antidepressant of another class; Stage III adds failure of an adequate tricyclic trial; Stage IV adds failure of an adequate monoamine oxidase inhibitor trial; and Stage V adds failure of a bilateral electroconvulsive therapy course. It is a ladder, and the ladder encodes a historical hierarchy in which the older and less tolerable agents sit higher up, which is both its clarity and its weakness: a patient given three modern agents in sequence cannot climb past the second rung. The procedure at the top rung is set out in the ECT and neurostimulation notes.

The **Maudsley Staging Method** is multidimensional. It is a points-based model combining three factors, treatment, severity of illness and duration of the presenting episode, scored so that the total runs from 3, marking minimal resistance, to 15, marking severe resistance. The published description gives the total range but not the point values for the three sub-scales, so only the total is printed here and the individual point allocations are to be taken from the original description. Learn the three dimensions, the direction of the scale and its two bounds. Because severity and chronicity carry points of their own, a patient with two failed trials and a severe, long-standing episode scores above a patient with two failed trials and a mild recent one, which is closer to how the difficulty is actually experienced in clinic.

How that severity dimension is anchored is stated in two places and the two must not be merged. The paper that first described the points-based model is the authority for the three factors and for the score range. The operationalised criteria attributed to the method in Paribello and colleagues' comparison table of staging systems are the ICD-10 severity categories, with validated symptom scales offered as an alternative, and that table credits its entry to a different publication on the same method rather than to the original description. Cite each to the paper that states it; a single citation covering both is wrong about one of them. The Massachusetts General Hospital model is named in the same reviews, but its scoring rules are not reproduced there, so no point values for it are given here.

6.7 Why none of this reaches the bedside

The candid position is that these staging models do not have proven predictive validity for treatment success and are not routinely used. That single sentence is worth more marks than any stage table, because it converts a list-learning answer into an appraisal. A staging system earns its

place by predicting something: which patient will respond to augmentation, which will need a procedural treatment, which will remit at all. These grade the past and have not been shown to forecast the future.

Formal consensus has said the same thing arithmetically. When the two leading staging tools were put to a Delphi vote, neither reached the 80 percent consensus threshold for recommendation: the DM-TRD scored 72 percent and the Maudsley staging model 45 percent. That same report states a different pair of figures for these two items in another passage of its own text; both cannot be the paper's figure, so only the pair given here is printed and the discrepancy is named rather than resolved.

The deeper trouble sits upstream of staging altogether. A systematic review by Brown and colleagues, published in the Canadian Journal of Psychiatry in 2019, identified 155 definitions of TRD in the published literature across 150 studies, and 48.4 percent of all definitions specified a requirement of at least 2 treatment failures. Both figures are that review's own count and are quoted on its authority. Read the second figure carefully, because it is the examinable one: a little under half of published definitions use the two-trial rule, which means that the population labelled resistant in one trial is not reliably the population labelled resistant in the next. Every efficacy figure quoted for a treatment in resistant depression inherits that imprecision, and it is the reason the switch, combination and augmentation evidence taken up later in this chapter has to be read against the definition each study used.

6.8 From resistant to difficult to treat

The terminology is moving, and the guidelines have moved ahead of the textbooks. The Maudsley Prescribing Guidelines record a move to characterise treatment-resistant depression as **difficult to treat depression**, on the basis that the former description implies that depression treatments are normally effective and that non-response is therefore somehow abnormal. That rationale is worth reproducing accurately, because it is an argument about base rates rather than about kindness of language: if a third of patients do not respond to two adequate trials, non-response is a common outcome of treatment and not a deviation from it.

Indian practice has already adopted the reframing in its own text. In the IPS CPG 2025 depression guideline, difficult-to-treat depression is used as the holistic framework, while treatment-resistant depression is retained when referring to medication options or to studies that use that definition. That sentence, rather than the two-trial rule, is the one that separates a current answer from a dated one: a candidate who offers only the failed-trial definition is answering the older guideline. The framework term governs the formulation and the narrower term governs the pharmacology, which is a division of labour rather than a change of mind, and the same logic explains why the resistant-schizophrenia material later in this chapter keeps its own vocabulary.

GOOD TO REMEMBER

The reframing changes the question asked at the review appointment. Resistance asks which agent has not yet been tried. Difficulty asks what is keeping this episode open, which returns the clinician to the five pre-label checks and to the contributors the reappraisal is meant to surface. Both questions are legitimate and only one of them can be answered by writing another prescription, so ask the second before the first.

Two failed adequate trials make the label; adequacy, adherence and a re-verified diagnosis are what make the label true.

7 · Switching, combination and augmentation strategies (lithium, atypical antipsychotics, thyroid)

A failed antidepressant trial forces a choice, not a new diagnosis. Once an adequate therapeutic dose has been given for an adequate duration, with optimisation considered against response and tolerability rather than pushed to the tolerability limit as a matter of routine, three structurally different manoeuvres are available: replace the drug, add a second antidepressant, or add something that is not an antidepressant at all. The marks sit on knowing that these are not interchangeable rungs of one ladder, that the order guidelines print them in is a sequence of decisions rather than a ranking the evidence supplies, and that almost every number attached to them belongs to one named body and is wrong under any other attribution. The definition and staging of resistance are taken up separately; what follows begins after resistance has been established.

7.1 The sequence, and what each rung actually is

The Indian 2025 guideline sets a six-step sequential ladder for **difficult-to-treat depression**, and it is the cleanest scaffold to answer from because it separates manoeuvres that an unstructured answer runs together.

- **Optimise** the current treatment: raise it to the maximum tolerated dose and monitor adverse reactions before concluding anything about response.
- **Switch** the antidepressant, choosing an agent with a different mechanism of action, such as an SSRI to an SNRI, or an SNRI to bupropion.
- **Augment**: add an agent that is not itself an antidepressant.
- **Combine**: add a second antidepressant from a different class, and the guideline attaches the words with caution to this rung specifically.
- **Neuromodulate**, with electroconvulsive therapy for severe and suicidal presentations; the modalities are taught in their own sections here.
- **Novel agents**, which is where the glutamatergic drugs sit; it has its own section here.

Two distinctions inside that list are examined more often than the list itself. **Augmentation** means adding an agent from outside the antidepressant classes, so the antidepressant continues

unchanged and the added agent works through a different mechanism. Combination means running two antidepressants together. The clinical consequence differs: an augmenting agent imports a new adverse-effect class and a new monitoring requirement, whereas a second antidepressant mostly intensifies pharmacology the patient already has, and with serotonergic pairs it raises the serotonin syndrome question taken up elsewhere here. The receptor pharmacology of each individual agent belongs to the Psychopharmacology I: principles and core drug classes notes.

MNEMONIC

Optimise, Switch, Augment, Combine, Neuromodulate, Novel

The six rungs in order. The first is the one candidates skip and the one that costs the mark: a trial at a subtherapeutic or otherwise inadequate dose has not failed, and the sequence does not begin until it has. Submaximal and subtherapeutic are not the same word. Many adequate trials run at a therapeutic dose below the maximum the patient could tolerate, so optimisation is a separate decision taken on response and tolerability, which is why the guideline gives it a rung of its own rather than folding it into the definition of failure.

7.2 Switching, and why its evidence is thinner than its position

Switching is the second rung in the Indian 2025 sequence and appears early in many algorithms, which invites the assumption that it is the best-supported second move. It is not. In STAR*D level 2, patients switched to another agent after citalopram failure reached HRSD-17 remission in similar proportions whichever agent they were given, with the difference between the three falling short of significance.

21.3%	17.6%	24.8%
SWITCH TO BUPROPION-SR	SWITCH TO SERTRALINE	SWITCH TO VENLAFAXINE-XR

The add-on arm of the same programme is examined nearly as often, and one point of nomenclature has to be settled before its numbers mean anything. STAR*D called both add-on strategies augmentation; under the taxonomy this chapter takes from the Indian guideline, adding bupropion is antidepressant combination and only adding buspirone is augmentation. Trivedi and colleagues, reporting that arm in the New England Journal of Medicine in 2006, record HRSD-17 remission of 29.7 percent with added sustained-release bupropion and 30.1 percent with added buspirone, and a lower dropout rate due to intolerance with bupropion, 12.5 percent against 20.6 percent. Both of those comparisons live inside the add-on cohort. That cohort was never randomised against the switch cohort, so no switch-versus-augmentation inference can be drawn from the two together: the numbers sit side by side because the programme ran both, not because anyone compared them, and an answer that reads the higher add-on figure as evidence that adding beats switching has overreached. The tolerability advantage is narrower still, belonging to bupropion against buspirone within the add-on comparison and saying nothing whatever about switching.

The harder counterweight comes from meta-analysis, which found no high-level evidence that switching the antidepressant is superior to simply **continuing the initial antidepressant**. That comparator is the one switching trials usually omit, and it matters because time alone produces response: a patient re-rated after a switch has had more treatment as well as a new drug, and designs that separate the two are rare.

RULE **The order of the ladder is a sequence of decisions, not a ranking of evidence.** The Indian ladder places switching early because it preserves monotherapy and avoids concurrent polypharmacy, not because it works better. That is not the same as costing nothing: switching carries its own discontinuation, interaction and new-drug adverse effects, so the trade is one adverse-effect burden for another rather than a burden avoided. Nor does the evidence supply a general ranking to put in the order's place. Comparative efficacy between switching, combining and augmenting is specific to the agents and the trial, and a candidate who can say why the order and the evidence part company, without inventing a hierarchy to replace the one they have just dismantled, has understood the topic.

7.3 Combination: which pairing earns the term

Combination carries a quantified benefit over monotherapy, and the population it was quantified in has to travel with it: the analysis is of acute depression, including but not confined to patients who had already failed a trial, so it does not establish combination as the best strategy in treatment-resistant depression specifically. A meta-analysis of 39 RCTs including 6,751 patients found combining antidepressants superior to **monotherapy**, with a standardised mean difference of 0.31. That benefit was unevenly spread across the pairings. Combinations built on a **presynaptic alpha-2 autoreceptor antagonist** were superior to other combinations, at 0.37. Bupropion combinations were not superior to monotherapy. A complementary autoreceptor mechanism is the proposed explanation for that split, and it is worth being able to state: blocking the presynaptic autoreceptor lifts the negative feedback that damps monoamine firing, which is a different lever from the one a reuptake inhibitor pulls. What the analysis actually demonstrates is an association by combination class. It does not test that mechanism, and the result is not a reason to say that bupropion merely duplicates the pharmacology of the drug it is added to. The same analysis carried an indication of publication bias, and an answer that quotes the effect size without that caveat is quoting half of it.

The Maudsley Prescribing Guidelines place the practical pairings in their tier of commonly used treatments, generally well supported by published literature. What that tier prints for treatment-resistant depression is a comparative list, not one regimen: three add-ons across two backbones, with no starting dose stated for any of them. An SSRI plus bupropion is entered as up to 400 mg/day, a ceiling and not a start, and the same evidence tier records that bupropion is not licensed for depression in the United Kingdom. Mianserin added to an SSRI or venlafaxine backbone in resistant depression is a flat figure, 30 mg/day, and the disadvantage recorded against it in the same evidence table is a risk of blood dyscrasia. Mirtazapine on those same backbones, again for resistant depression, is 30 to 45 mg/day, and one randomised trial showed no advantage for mirtazapine added to an SSRI or SNRI. Both autoreceptor pairings carry a theoretical risk of serotonin syndrome that the patient is to be informed of.

7.4 Lithium augmentation, and the number no Indian source states

Lithium is the drug primarily used as an **adjunct** in treatment-resistant depression according to the Indian 2017 guideline, which reports it useful in over 50 percent of antidepressant non-responders and usually well tolerated. Its most examinable feature is the latency: the interval before full response to adjunctive lithium is reported as **several days to 3 weeks**. That figure decides when a trial of the addition may fairly be called a failure, and it is strikingly short beside the antidepressant trial that preceded it.

No serum lithium target for antidepressant augmentation is carried by the Indian guidance these notes rest on, so the bands below belong to the bodies that state them. For lithium augmentation of an antidepressant the Maudsley target is a plasma level of **0.4 to 0.8 mmol/L** initially, increasing to up to 1.0 mmol/L if the response is suboptimal, and the same entry records plasma level monitoring as essential. NICE, for lithium used to augment antidepressant medication, states that plasma lithium levels should not exceed 1.0 mmol/L, that therapeutic levels for that purpose are usually at or above 0.4 mmol/L, and that levels of 0.4 to 0.6 mmol/L should be considered for older people aged 65 or above. The two are not the same kind of statement. The Maudsley gives a band to aim at. NICE gives a ceiling, and its lower figure is a usual practice rather than a floor: augmentation levels are usually, not invariably, at or above 0.4 mmol/L, so a level below it is not thereby declared subtherapeutic.

TRAP **Putting an Indian attribution on a lithium augmentation target.** The Indian guidance names lithium as the adjunct of first choice and states no serum level for that indication, so a level quoted under an Indian guideline's name is a fabricated citation even when the digits happen to be right. This is a different failure from quoting the wrong figure, and harder to spot: the number survives a memory check while the source does not. Quote the body that states the level, by name, in the same sentence as the level.

IN INDIA

The Indian guidance is the better scaffold and the thinner source of numbers, and the practical decision follows from that split. Take the sequence from the Indian ladder, which puts optimisation before switching and names lithium, thyroid hormone or an atypical antipsychotic as the augmenting options with the greatest evidence for lithium and aripiprazole; take the level, the dose and the monitoring interval from whichever body states them, and name that body in the same breath. Among the Indian figures in this section, the short latency reported for adjunctive lithium is the one that changes bedside behaviour most directly, because it fixes how long to hold the addition before judging it.

7.5 Thyroid hormone: one agent, two doses, two bodies

Tri-iodothyronine is the preparation used for this purpose rather than levothyroxine, and it is given to patients who are already **euthyroid**, which is what makes it counter-intuitive at the bedside and durable in vivas. The rationale is not replacement of a deficiency but potentiation of an incomplete antidepressant response, which is why thyroid function is followed as a safety measure rather than as the endpoint being treated. The Maudsley files it a tier below lithium,

aripiprazole and quetiapine, among less commonly used treatments that are nonetheless well supported by published evidence.

The two doses in circulation are genuinely different instructions from different bodies. The Indian 2017 guideline gives thyroid hormone augmentation of an antidepressant, even in euthyroid patients, at **25 micrograms/day** of tri-iodothyronine increased to 50 micrograms/day after a week. The Maudsley guidelines give tri-iodothyronine augmentation of an antidepressant at 20 to 50 micrograms/day, note that higher doses have been used safely, and require clinical and biochemical thyroid function monitoring alongside. The upper figures coincide; the starting point and the shape of the instruction do not. Neither body states a maximum for tri-iodothyronine. The Indian guideline gives a start and a step and stops there, with no ceiling named above either. The Maudsley tri-iodothyronine entry travels in the opposite direction, recording that higher doses have been used safely, which is the reverse of naming a limit, and it gives no starting dose at all, only the band inside which the prescriber chooses. So the figure the two documents share is where two instructions happen to meet, and not a ceiling either of them has set.

TRAP **Printing one tri-iodothyronine dose as the tri-iodothyronine dose.** The Indian figure is a titration with a fixed start and a fixed step at a fixed interval; the United Kingdom figure is a range within which the clinician chooses, with an explicit note that higher doses have been used safely. Writing only the range drops the Indian instruction; writing only the start and the step drops the range and that note. Write both, name the body against each, and never average them into a single band that neither source states.

7.6 Atypical antipsychotic augmentation: agent, dose and the trade it buys

This is the strategy modern guidelines lean on hardest. **Atypical antipsychotic** augmentation is named in the Indian ladder beside lithium and thyroid hormone, with the greatest evidence available for lithium and aripiprazole. The Indian text names the agent and states no dose, so the doses that follow are the Maudsley's.

Aripiprazole augmentation of an antidepressant is used at **2 to 20 mg/day**, with low doses of 2 to 10 mg/day potentially effective, and akathisia and restlessness common at standard doses of 10 mg/day or more, insomnia being the other effect that may become problematic. Quetiapine augmentation of a depressive illness is used at 150 mg or 300 mg a day added to an SSRI or SNRI, an option the Maudsley credits with a good evidence base and usual good tolerability, set against dry mouth, sedation and constipation as the problematic effects and weight gain as the longer-term risk. The olanzapine and fluoxetine combination for depression is used at 6.25 to 12.5 mg olanzapine plus 25 to 50 mg fluoxetine daily, which is the United States licensed dose; where combination formulations are not available, the Maudsley footnote says 5 mg plus 20 mg rising to 10 mg plus 40 mg seems reasonable, which is the line that matters wherever the fixed-dose product is absent.

None of those three entries is a complete prescribing instruction, and it is worth saying so beside the figures rather than leaving a reader to notice. The aripiprazole entry records the effective low

range and the dose at which akathisia becomes common, and no dose to begin at. The quetiapine entry records the two amounts used in augmentation, with no starting dose below them and no maximum above them. The olanzapine and fluoxetine entry records a licensed range and an alternative for units without the fixed-dose product, with no starting dose for the licensed range and no ceiling stated for that combination at all; the alternative does carry a start, since it is written as a titration. Its upper pair is what the licence permits, which is not the same statement as a limit the guideline has set. Read all three as the amounts these table rows record, and no further.

CANMAT 2023 sharpens the choice inside the class, listing aripiprazole and brexpiprazole as first-line adjuncts for difficult-to-treat depression on efficacy and tolerability, with cariprazine second-line because there are fewer studies of it. Whether brexpiprazole and cariprazine are approved and marketed here, and at which strengths, is not something the sources behind these notes establish, and the same is true of the marketed strength of tri-iodothyronine.

GOOD TO REMEMBER

Start aripiprazole augmentation at the bottom of its range and give it time before moving up. Akathisia and restlessness belong to the standard doses rather than the low ones, so a patient who becomes agitated and unable to settle a few days after the addition is more likely to be akathisic than more depressed. The reflex is to read the agitation as worsening illness and push the dose, which makes both problems worse.

7.7 Reading the guideline bodies against each other

Four bodies are quotable on augmentation and they do not carry the same agents at the same strength. BAP 2015 grades first-line augmentation as quetiapine, aripiprazole or lithium, each at grade A, and second-line as risperidone at grade A with olanzapine, tri-iodothyronine and mirtazapine at grade B. The caveat printed with those grades is the teaching point and is regularly dropped: the evidence derives mainly from studies in which lithium and tri-iodothyronine were added to tricyclics, and the other drugs were added to SSRIs or SNRIs. A grade A for lithium is therefore a grade A earned in a prescribing context that has largely gone.

AGENT	POSITION IN THE GRADED BRITISH LIST	NAMED IN THE INDIAN LADDER	WHERE A PRINTABLE DOSE OR LEVEL COMES FROM
Lithium	First-line, grade A	Named, and one of the two agents of greatest evidence	The Maudsley and NICE, as a plasma level rather than a dose
Aripiprazole	First-line, grade A	Named, and the other agent of greatest evidence	The Maudsley
Quetiapine	First-line, grade A	Covered by the atypical antipsychotic option, not named	The Maudsley

AGENT	POSITION IN THE GRADED BRITISH LIST	NAMED IN THE INDIAN LADDER	WHERE A PRINTABLE DOSE OR LEVEL COMES FROM
Risperidone	Second-line, grade A	Covered by the atypical antipsychotic option, not named	No dose is stated
Olanzapine	Second-line, grade B	Covered by the atypical antipsychotic option, not named	The Maudsley, only in combination with fluoxetine
Tri-iodothyronine	Second-line, grade B	Named, as thyroid hormone	The earlier Indian guideline and the Maudsley, and they differ
Mirtazapine	Second-line, grade B	Not an augmenting agent there; it is a combination partner	The Maudsley
Brexpiprazole	Not among the graded agents	Not named	The Canadian list, as a first-line adjunct, with no dose stated
Cariprazine	Not among the graded agents	Not named	The Canadian list, as a second-line adjunct, with no dose stated

Read down the third column and the Indian position becomes legible: it names a class where the other lists name molecules, a deliberate looseness written to survive a formulary that varies between institutions. The cost is that the Indian document cannot be quoted for a dose.

7.8 Monitoring what you have added

Every augmenting agent brings its own surveillance, and the schedules are the part of this topic a written answer most often omits. For lithium used to augment antidepressant medication, NICE assesses weight, renal and thyroid function and calcium levels before treatment and then monitors at least every 6 months during treatment. CANMAT 2023 states that when adjunctive agents are used, laboratory tests tailored to the specific medication should be monitored, and gives as its worked example that patients taking lithium should have lithium levels, electrolytes, calcium, creatinine, eGFR and TSH checked at baseline and every 6 to 12 months thereafter, or whenever the clinical status or the lithium dose changes. Quote that interval with a caveat attached. Neither this monitoring sentence nor the adjunct ranking taken from the same update could be located in the copy used for these notes, so the six to twelve month figure is unverified here and the NICE schedule is the one to rely on until it is confirmed.

Those are two different instructions from two bodies and neither is offered here as the resolution of the other. They also do not cover quite the same thing, and that is the part most easily mishandled. NICE sets a separate and more frequent schedule for the serum lithium levels themselves during the first year of treatment; the intervals in that schedule are not carried by these notes, and it would understate NICE to present its six-monthly panel as its instruction for

lithium levels in the first year. Say the two bodies disagree, give each figure with its body, and be explicit that the British level-monitoring schedule is tighter than the panel interval implies.

For an antipsychotic used to treat depression, NICE directs that monitoring is carried out as indicated in the summary of product characteristics for the individual medicine, and says this may include monitoring weight weekly for the first 6 weeks, then at 12 weeks, at 1 year and annually, and monitoring fasting blood glucose or HbA1c and fasting lipids at 12 weeks, at 1 year and then annually. Read the modal exactly. The schedule is offered as what product-specific monitoring may include, not as a schedule the guideline itself specifies, and an answer that presents it as a fixed requirement has silently upgraded the strength of the recommendation. Tri-iodothyronine augmentation carries its own requirement, since the Maudsley requires clinical and biochemical thyroid function monitoring for it. The wider cardiometabolic and QT surveillance that spans the psychotropics is taught in its own section here.

PITFALL

Adding quetiapine or olanzapine to an antidepressant and filing the metabolic monitoring under something that only applies to psychosis. The augmenting antipsychotic is prescribed by the clinician who has been treating a depressive illness, often in a setting that has never taken a baseline weight or a fasting glucose, and the weight gain listed as the longer-term risk accrues while nobody is measuring it. The companion instruction is to check baseline pulse and blood pressure, weight, nutritional status, diet, level of physical activity, fasting blood glucose or HbA1c and fasting lipids before the antipsychotic is started, not at the first review.

Optimise before you move: switching preserves monotherapy but brings its own discontinuation and new-drug risks and has no high-level evidence of superiority over simply continuing, the comparison between the three manoeuvres is specific to the agents and the trial rather than a general ranking, and every level, dose and monitoring interval you quote belongs to a named body and must travel with its name.

8 · Ketamine and esketamine in treatment-resistant depression (cross-reference to recent advances)

Glutamate arrived late in the antidepressant story and changed the timescale. Ketamine and its S-enantiomer esketamine are the worked example of what a resistant depressive episode does to prescribing. Here is an agent whose effect is measured in hours rather than in weeks of waiting, whose route decides which evidence applies to it, whose durability is quantified for one formulation and merely described for the other, and whose delivery needs a monitored chair rather than a prescription pad. The marks sit on route, eligibility, monitoring and durability, and candidates lose them by treating ketamine and esketamine as one drug with one dose. The receptor pharmacology of glutamatergic agents as a drug class belongs to the Psychopharmacology I: principles and core drug classes notes; what follows is the decision that has to be made once two adequate antidepressant trials have failed.

8.1 Where the glutamatergic agents enter, and who is eligible

The 2025 Indian Psychiatric Society depression guideline puts novel agents, naming ketamine and esketamine alongside dextromethorphan-bupropion, in the last row of its sequential intervention table for **difficult-to-treat depression**, qualified "wherever available" and annotated "Limited long-term data" in the notes column. The position in the table is itself the teaching. Within that Indian table these agents come after the conventional optimisation, switching, augmentation and combination strategies and after neuromodulation, each of which is a topic in its own right; and the recommendation is conditioned on supply rather than asserting that the agents can be obtained. The ordering belongs to this guideline rather than to the field: the Canadian guidance in the next paragraph puts the same agents earlier, so an answer should name the body whose sequence it is describing.

The Canadian guidance shows how quickly the position moved, and the direction of travel is examinable. CANMAT in 2016 considered ketamine an experimental treatment and recommended that its use be limited to academic depression treatment centres, its adjunctive-agent table listing ketamine as Experimental, Level 1 evidence, at 0.5 mg/kg as a single intravenous dose for acute treatment. By 2023 the same body recommended adjunctive ketamine and intranasal esketamine at second line, which is an advance on the experimental status of 2016 and not a retreat from it. CANMAT's own verb is downgraded, but what it downgrades is the placement the efficacy evidence alone would have supported, and the stated reason is worth memorising because it is not loss of efficacy: side effect potential and feasibility issues, specifically the required monitoring of blood pressure and dissociative side effects. The same update pushed the alternate routes of racemic ketamine, oral, nasal and intramuscular, down to third-line, given highly variable protocols and mixed results. That the second-line agent is the intravenous racemic route is a reading derived from the surrounding paragraph rather than a phrase the guideline itself prints.

Eligibility is defined most precisely not by a guideline but by a licence. The Summary of Product Characteristics for esketamine nasal spray indicates it, in combination with an SSRI or SNRI, for adults with treatment-resistant major depressive disorder who have not responded to at least two different antidepressant treatments in the current moderate to severe depressive episode. Read the three constraints together: the failures must be to antidepressant treatments, they must lie inside the current episode, and that episode must be moderate to severe. This is the eligibility criterion written into one product's licence, and it is not the same thing as the general definition and staging of treatment resistance, which is a separate construct.

8.2 Mechanism, and the shape of the acute effect

Ketamine is an antagonist at the NMDA glutamate receptor, and esketamine is its S-enantiomer formulated as a nasal spray. The prevailing account of the antidepressant action is not the block itself but what follows it: preferential action on inhibitory interneurons is thought to disinhibit pyramidal output, producing a brief rise in glutamate that acts at AMPA receptors and drives a rapid strengthening of synaptic connections in prefrontal circuits. That account is a hypothesis

under active revision, not a settled mechanism, and an answer that states it as established has overreached. What is not in doubt is the clinical shape of the effect.

The Maudsley Prescribing Guidelines, describing the first controlled trial, record that a single **subanaesthetic** intravenous infusion of ketamine at 0.5 mg/kg produced an antidepressant effect within hours, and that the effect increased progressively up to 3 days after administration. Two consequences follow for practice. The peak does not fall at the end of the infusion, so a patient written off as a non-responder in the recovery chair may have been judged too early; and a benefit that builds over days is a different phenomenon from the dissociative experience of the infusion hour, which is why examiners like the distinction between the psychotomimetic and the therapeutic effect.

RULE **Ketamine and esketamine are not one treatment, and the route is what decides the evidence.** Molecule, enantiomer, route, licence and monitoring requirement travel together as a single package. A statement about an intravenous racemic infusion does not transfer to a nasal spray, and the guideline bodies grade the routes separately for precisely that reason.

8.3 Route and formulation: two standard sources, one drug, different numbers

Route is where most of the examinable detail sits, and it is also where the two texts on an Indian shelf disagree without either being wrong. The rule for an answer is to print both, attributed, and never to split the difference. Note as you read the table that the Indian document supplies figures for routes it simultaneously declines to endorse, and that **racemic ketamine** and esketamine are graded separately at every step.

ROUTE	INDIAN PSYCHIATRIC SOCIETY DEPRESSION GUIDELINE, 2025	MAUDSLEY PRESCRIBING GUIDELINES, TREATMENT-RESISTANT DEPRESSION
Intravenous racemic ketamine	0.5 to 1.0 mg/kg, infused over 40 to 60 minutes, twice weekly for 2 weeks in the acute course	**0.5 mg/kg over 40 minutes**, named the gold standard for off-label ketamine administration, with the best supporting evidence for efficacy, and possibly more effective than intranasal administration
Oral ketamine	2 to 3 mg/kg, mixed with around 100 mL of purified water and sipped over 10 to 15 minutes for the first session, with the available evidence called preliminary and currently insufficient to support routine clinical use	0.5 to 1.25 mg/kg, effective but with no licensed products
Sublingual ketamine	Not established, with a suggested start at 10 mg placed under the tongue and held for 2 minutes without swallowing	Oral or sublingual doses considered between intravenous treatments
Intranasal esketamine	56 to 84 mg twice weekly for 4 weeks of induction	28 to 84 mg, licensed in most countries, where the oral evidence above has no licensed product
Intranasal racemic ketamine	50 to 150 mg twice weekly	No separate figure in the treatment-resistant depression table

Two disagreements inside that table are examinable in their own right. The first is infusion duration: the Indian range is wider at its upper end than the single British figure, and the British figure is also the duration used in the controlled trial that text describes. The range contains the point, so the two are compatible at the bedside and not interchangeable in an answer. Name the body you are quoting and give its own figure.

Something else is true of every cell in that table, and it is easy to miss because the cells look finished. No cell in it states a maximum. For ketamine, intravenous, oral or sublingual, and for esketamine by the intranasal route, the Indian guideline gives an amount and either a vehicle or an infusion time, and then stops, with no ceiling named for any of them. Nor does every row carry a frequency, and the rows differ on that point in a way the layout hides. The intravenous and intranasal rows state twice weekly. The oral row describes the first session only, with no interval between sessions and no course length. The sublingual row states neither, and its two minutes is how long the dose is held under the tongue rather than how often it is given, which is the likeliest cell here to be misread as a schedule. The Maudsley gives the intravenous ketamine figure it calls the gold standard and states no limit above it, and no interval between infusions either. Only the sublingual ketamine row is written as somewhere to start, and its own source calls that route not established while suggesting the figure; for every other ketamine and esketamine cell here there is no starting dose, because the figure recorded is the amount to give

rather than the first step of a titration. So the number a reader meets in any of these cells is what the source recorded for that route, and not the top of a permitted range.

TRAP **Averaging the two published oral ketamine doses.** For depression the Maudsley gives oral ketamine at 0.5 to 1.25 mg/kg, calling it effective on that evidence while recording that no licensed oral product exists, and the Indian guideline gives 2 to 3 mg/kg, a several-fold separation for one route in one indication. Neither is a misprint and no average exists between them: they rest on different literatures, and the Indian guideline attaches to its own figure the statement that the evidence for oral ketamine is preliminary and currently insufficient to support routine clinical use. A candidate who writes a single oral range has silently chosen a source and hidden the choice.

8.4 The esketamine schedule, and the second licensed indication

Esketamine is the one formulation with a fixed published schedule rather than a range, and it is the one most often reproduced wrongly, because the frequency steps down twice. In adults under 65 with treatment-resistant depression the label sets **56 mg on the first day, then 56 mg or 84 mg twice a week through the first four weeks**, the same dose once weekly through the fifth to eighth weeks, and from the ninth week onward the same dose either once every two weeks or once weekly. Three phases, two step-downs, one dose pair throughout: that is the whole table, and it is easier held as a shape than as a grid. The esketamine label sets no maximum above that pair. Its higher figure is a maintenance dose that recurs at every phase, so treating it as the licensed ceiling reads a schedule as a limit, which the schedule does not say.

The label attaches a durability instruction to that schedule. After depressive symptoms improve, treatment is recommended for **at least 6 months**. It is the instruction most often dropped by an answer that has carefully memorised the induction phase and stopped there.

There is a second licence, and it is not for resistance. Esketamine is also licensed for the acute short-term treatment of depressive symptoms constituting a psychiatric emergency, co-administered with oral antidepressant therapy, and for that indication the recommended dosage in adults under 65 is 84 mg, the maximum the label recommends for this indication, given twice per week for 4 weeks, with reduction to 56 mg made on tolerability. The two indications differ in more than the number. The emergency indication opens at the higher of the two doses and comes down if the patient cannot tolerate it, whereas the resistance indication opens lower and may go up; a stem that gives you the entry dose has already told you which licence it is testing. Note the shape of the emergency esketamine entry too. It records the dose for that indication and no starting dose beneath it, so the reduction it describes is a response to tolerability and not the second rung of a titration.

8.5 Durability: quantified for one formulation, described for the other

This is the discriminating section of the topic, and the asymmetry in it is uncomfortable: the strong maintenance evidence belongs to intranasal esketamine, whose Indian availability is unverified here, while the intravenous racemic infusion has no equivalent relapse-prevention trial. Which of the two a given service can actually deliver is a local question, to be established rather than assumed.

For racemic ketamine the Indian guideline describes an acute course of six to eight infusions, given twice or three times weekly, and states that continuation and maintenance ketamine may be given at an individualised frequency but that most of the evidence along those lines is anecdotal. Its own route table, printed above, describes that acute course differently, as twice weekly for 2 weeks. The two do not reconcile, and both are in the same document, so the honest answer names both rather than choosing between them. The Maudsley schedule is compatible with that and equally unquantified: an **induction phase** of once or twice a week, then a maintenance phase, according to response, of weekly, then every 2 weeks or even monthly, with oral or sublingual doses considered between the infusions.

Esketamine has the evidence the racemic infusion lacks. In SUSTAIN-1, a **randomised withdrawal** trial, continuing esketamine nasal spray with an oral antidepressant reduced relapse risk by 51 percent in patients who had achieved stable remission and by 70 percent in those who had achieved stable response. The design is the point as much as the result: patients were stabilised first and then randomised to continue or to stop, so what the trial measures is the value of not stopping.

0.49	6	0.30	4
HAZARD RATIO, STABLE REMITTERS	NUMBER NEEDED TO TREAT, REMITTERS	HAZARD RATIO, STABLE RESPONDERS	NUMBER NEEDED TO TREAT, RESPONDERS

The head-to-head comparison points the same way. In ESCAPE-TRD, esketamine nasal spray was superior to extended-release quetiapine, each added to an SSRI or SNRI, for remission at 8 weeks, reached by 27.1 percent against 17.6 percent, and for remission at that point sustained without relapse through 32 weeks, achieved by 21.7 percent against 14.1 percent. Two design features reported with the trial should travel with those figures whenever they are quoted: the comparison was open-label with the raters kept unaware of allocation, and patients who discontinued were counted as having an unfavourable outcome.

TRAP **Carrying the esketamine relapse figures across to racemic ketamine.** The hazard ratios and numbers needed to treat above belong to intranasal esketamine continued alongside an oral antidepressant in a randomised withdrawal design. No comparable maintenance trial of the intravenous racemic infusion with a relapse endpoint is available to this account, and the Indian guideline itself calls the continuation evidence anecdotal. The defensible sentence is that durability is quantified for esketamine and described only qualitatively for racemic ketamine, and a candidate who transfers a hazard ratio between them has invented a result.

8.6 What has to be in place before, during and after a session

This is where the modality stops being a prescription and becomes a service, and it is the part an answer built only from dose tables leaves out. Before administration the Maudsley recommends a physical examination including baseline blood pressure, full blood count, liver function test, thyroid function test and an electrocardiogram, with blood pressure and heart rate monitored during and after administration.

During and after the session the two standard sources set different clocks, and both are examinable. The Indian guideline describes continuous monitoring of consciousness, heart rate,

blood pressure, respiratory rate and oxygen saturation at baseline and periodically, offering every 15 minutes as its own example, throughout the session and continuing for up to **2 to 3 hours** after the session ends. The Maudsley, for racemic ketamine, has the patient observed by a trained clinician during dosing and for 1 hour after administration. For an Indian examination the Indian figure governs and the British one is quoted as the comparator; the two were set by different bodies and must not be averaged into a third period that neither recommends.

> **PITFALL**
>
> Treating the observation period as a courtesy rather than a competence requirement. Ketamine given at antidepressant doses causes significant **dissociative symptoms**, so patients must be observed by a trained clinician, and because **laryngospasm** has been reported the observing clinician should be trained in intermediate or advanced life support. Being able to obtain the drug is therefore not the same thing as being able to give it.

Esketamine carries its own haemodynamic clock. Its label records transient rises in systolic and diastolic blood pressure that peak at approximately 40 minutes after administration and last approximately 1 to 2 hours, and it times the post-dose reassessment of blood pressure to that same 40 minute point. That single figure explains the shape of the whole session: the reassessment is set against the expected peak, not against the end of the dose. For patients kept on longer-term ketamine or esketamine treatment the Canadian guidance adds periodic urinalysis to check for cystitis, a potential long-term adverse effect of ketamine use, which is the one piece of monitoring that outlives the session itself.

8.7 The Indian position: a controlled drug and a cheap route the guideline will not endorse

Ketamine is a **controlled drug** in India, so the prescribing, storage and record-keeping frame is statutory before it is clinical, and that is a fact about the drug rather than about any one hospital. Which instrument and which schedule govern it is a separate question, and the sources read here do not settle it, so this account prints no schedule number. The one Indian paper read here calls ketamine a Schedule III controlled substance under the Narcotic Drugs and Psychotropic Substances Act, 1985, but that is the schedule language of United States controlled-substances law rather than a schedule of the Indian Act; beside it sit a United Nations country listing recording ketamine's inclusion among the psychotropic substances under that Act, and secondary sources placing ketamine hydrochloride in Schedule X of the Drugs and Cosmetics Rules. The Gazette notification that would decide between them could not be consulted for this chapter. Carry the statutory obligation into an answer and leave the schedule number out of it.

Set against that is what one Indian source reports. A 2024 correspondence article in an Indian journal states that the administration of oral ketamine is prevalent in India because it is inexpensive and easy to administer in facilities with limited access to healthcare resources. The attribution has to stay visible, because that sentence is a correspondence article rather than a prevalence study, and no Indian prevalence figure or national register of ketamine use was found. The tension is still the examinable point: the cheapest and most deliverable route is precisely the

route the Indian guideline prints a dose for and then declines to endorse for routine use, and the route the Canadian guidance places third-line.

On availability the guideline is careful and an answer should be too. It lists these agents as novel agents "wherever available", which conditions the recommendation on supply rather than asserting it. In the same passage the same authors do assert Indian availability for something else, recording that **dextromethorphan-bupropion**, approved in the USA in 2022 for major depressive disorder, is now available in India. The contrast is deliberate and it is usable in an answer: these authors state availability when they mean it, and they did not state it for esketamine. No Indian regulatory or market status for intranasal esketamine is asserted by any source used here, and the correct examination answer names that absence rather than filling it.

IN INDIA

Where a unit can obtain the drug at all, the decision the sources support is to run the intravenous route in a setting that can meet the guideline's own monitoring description, and not to substitute the oral route merely because it is simpler: the Indian guideline prints an oral dose and in the same table refuses to endorse it for routine use. The pre-treatment workup the British text asks for should be in hand before the first dose wherever a service offers ketamine, and the local availability of those five investigations, along with the referral route for any the unit cannot do itself, is to be confirmed rather than assumed. And because the drug is controlled in India whichever instrument turns out to govern it, the record-keeping is part of the treatment plan rather than an afterthought to it.

8.8 Choosing between them, and what a complete answer contains

Reduced to four decisions, the topic becomes answerable in whatever form the question takes.

- Route decides the evidence: the intravenous racemic infusion carries the best efficacy data, the nasal spray the licence, the fixed schedule and the maintenance trial.
- Eligibility is drawn from a licence rather than an impression, and the licence counts failed antidepressant treatments inside the current moderate to severe episode.
- Monitoring, not the drug, is the rate-limiting step: a trained observer, a defined observation period, and for esketamine a reassessment timed to the pressure peak.
- Durability is a separate question from response, and it has a different quality of answer for each formulation.

GOOD TO REMEMBER

The commonest real-world error here is not a dosing error. It is offering the treatment in a room that cannot observe the patient afterwards. Where the post-dose observation period cannot be staffed, the honest position is that the service is not yet able to offer the treatment, rather than that the treatment is unsuitable for the patient.

MNEMONIC

Route, Evidence, Observation, Durability

The four axes on which every question about these agents separates: which route is asked about, whose evidence supports it, how long the patient is watched afterwards, and what happens to the gain once the acute course ends. A durability claim made without naming the formulation is a wrong one.

Route decides the evidence, the licence decides eligibility, and the observation period decides whether the treatment can be offered at all.

9 · Treatment-resistant schizophrenia and the role of clozapine

Treatment resistance is a verdict on the prescribing before it is a verdict on the patient. Almost all the marks in this topic sit on the entry condition rather than on what follows it: what a guideline demands before the label may be used at all, what makes a failed trial a genuine failure, and which of the ordinary explanations for apparent non-response has been left unexcluded. Once resistance is properly established the drug choice narrows sharply, and the examinable content shifts to exposure, to dose and to blood counts. The schizophrenia management algorithm itself, and the order in which agents are tried, belong to the Schizophrenia: management, antipsychotics and adverse effects notes; what follows here is the pharmacological reasoning that sits underneath that algorithm and decides whether it should ever have been entered.

9.1 What the label requires before it may be applied

The current Indian definition of **treatment-resistant schizophrenia** is a compound one, and every limb of it is doing work. The 2025 Indian Psychiatric Society schizophrenia guideline requires failure across **at least 2 adequate trials** of different antipsychotics; it requires **confirmed medication adherence**, established rather than assumed; and it requires that moderate or severe symptoms and psychosocial dysfunction both persist despite that treatment. The comparator is easy to lose in an answer. A floor of two is not a rule of two, and a candidate who writes that resistance is declared after exactly two trials has quietly converted a minimum into a maximum.

The adherence limb is the one most often skipped, and it is the one that changes what happens next. Adherence is not a historical detail collected for completeness; it is written into the definition as a precondition, so an unconfirmed adherence history means the definition has not been satisfied and the label cannot yet be applied. The third limb keeps the judgement clinical rather than symptomatic: residual symptoms do not qualify a patient whose function has recovered, because the guideline binds severity and functional cost together rather than accepting either alone.

RULE **Resistance is a conjunction, never a count.** A tally of failed trials establishes nothing by itself. The trials must have been adequate, the adherence must have been confirmed rather than presumed, and the residual illness must be both severe and functionally costly. Fail any one of those and what is in front of you is inadequately treated illness, which is a different clinical problem with a different answer.

9.2 What makes a trial adequate: dose, duration and the drug actually taken

An **adequate antipsychotic trial** has a floor written into Indian guidance rather than left to judgement. The earlier 2017 Indian Psychiatric Society schizophrenia guideline set the minimum for all antipsychotics as the highest tolerable dose sustained for **6 to 8 weeks**, and carved out one exception: for clozapine the minimum period of treatment is **at least 3 to 6 months**. Two variables are specified and both must be satisfied before a trial counts: the dose pushed to what the patient could actually tolerate, and that dose held long enough for a response to declare itself.

This is where most apparent resistance dissolves on questioning. A trial fails the definition, rather than the drug failing the patient, in four recognisable ways.

- The dose never reached the highest tolerable level, usually because a tolerable adverse effect was read as a ceiling and the titration stopped there.
- The duration fell short, the agent being switched at the point where the prescriber lost patience rather than at the point where the trial actually ended.
- The drug was not taken, in which case no trial occurred at any dose and over any interval, and the correct next step is not a new molecule.
- The drug was taken but the patient never achieved the exposure the dose implied, which looks identical to non-adherence at the bedside and separates from it only on measurement.

A failed trial belongs in the tally the definition asks for once adequate dose, adequate duration and confirmed adherence have been established. The fourth is not a precondition on every antipsychotic trial: it is what a measured level investigates, where monitoring is indicated and interpretable, chiefly clozapine. The clozapine exception is asymmetrical by design. The minimum for this one drug is several times the minimum that satisfies every other antipsychotic, and the guideline records it as an exception rather than as a preference.

TRAP **Applying the standard adequate-trial yardstick to clozapine.** A patient taken off clozapine at the point where any other antipsychotic would fairly have been abandoned has not had a clozapine trial at all and cannot be described as having failed one. The trap is attractive because the yardstick is correct for every other agent, and because notes recording the switch rarely record how long the drug ran.

9.3 Pseudoresistance, and what a measured level settles

A level converts a clinical impression into a measurement. Where a patient on clozapine fails to improve, the Indian Psychiatric Society schizophrenia guideline names **inadequate response** among the conditions under which **therapeutic drug monitoring** of clozapine is usually recommended, and it states what the measurement is for: to establish whether under-exposure

explains the failure before the patient is called clozapine-resistant. Read the modal as it is written. Usually recommended is not required, and an answer that upgrades the strength of a recommendation has misreported the guideline as surely as one that misquotes a number.

The reasoning this licenses is the pharmacological lens on this topic, and it is reasoning about exposure rather than a verdict on treatment. A concentration below the accepted band raises under-exposure, through under-dosing, rapid metabolism, enzyme induction or partial adherence, but the sampling time, the steady state, the adherence history, smoking, interacting drugs and any intercurrent inflammatory illness are checked before the dose is touched, and dose adjustment is one response among several. A concentration within the band documents exposure at the moment of sampling; adequacy still needs the duration, the adherence history and the clinical criteria before the failure can be attributed to the drug. A concentration at or near zero in a patient who is collecting prescriptions strongly raises covert non-adherence or interrupted exposure, once the timing and the validity of the assay have been checked. These readings can coexist, so the level narrows the differential and never settles it alone; it is read alongside the clinical assessment, not instead of it. The threshold itself, the sampling technique and the intervals that surround it are taught in the therapeutic drug monitoring material and are not restated here.

PITFALL

Writing "failed clozapine" into a discharge summary without recording the dose reached, the duration sustained, or whether exposure was ever measured. The next prescriber inherits a closed door and no way to test whether it was ever locked. The guideline's sequence, measurement before label, cannot be run retrospectively on a note that omits the exposure.

9.4 Where clozapine enters, and how the entry condition changed

Indian guidance has stated the clozapine entry point twice, and the difference between the two statements is examinable in its own right. The earlier guideline held that clozapine need to be considered after the failure of sequential trials with 2 antipsychotics, at least one of them a second-generation agent. The current definition of resistance drops the class requirement altogether: what it asks for is failure across adequate trials of different antipsychotics with adherence confirmed, and it names no class anywhere. Reciting the second-generation requirement as current Indian practice is therefore an error of vintage rather than of fact, and a well-read candidate is likelier to make it than a poorly read one.

The modal in the earlier wording repays attention too. Clozapine need to be considered is an instruction to consider, which is what a guideline says when the decision is still the clinician's. The current guideline is stronger: its recommendation tables direct a switch to clozapine after at least two failed adequate trials, a direction that contraindications, tolerability and the patient's own decision still qualify. The pharmacological argument for moving early rests on the alternatives rather than on the guideline strength: once two adequate trials of different agents have failed, the probability that a third agent from the same broad pharmacological space will succeed is not what drives the decision, and clozapine is the option with a distinct claim on that

ground. The evidence-ranked ladder of what may be added to or substituted for an antipsychotic, and the switch-versus-combine-versus-augment decision itself, are taught under the switching, combination and augmentation strategies.

9.5 Dose is not exposure, and the Indian patient needs less of it

The dose written on the card and the concentration the patient achieves are different quantities, and the gap between them is populated by identifiable variables rather than by noise. The Maudsley Prescribing Guidelines record that clozapine plasma levels run lower in younger patients, in males and in smokers, and higher in Asians, and they state the practical consequence directly: markedly lower clozapine doses are required in East Asians, Indians and Bangladeshis. That is a statement about the patients in front of this reader, not a footnote about somebody else's population. The kinetic principles that generate the variability, and the cytochrome map that underlies most of it, belong to the Psychopharmacology I: principles and core drug classes notes.

No Indian milligram figure follows from that observation, and none should be manufactured to fill the space; the source gives a direction of effect and nothing more. The nearest dose anchor available is regional rather than Indian. The Malaysian Ministry of Health schizophrenia guideline starts clozapine at 12.5 mg/day, gives a target dose of 300 to 900 mg/day in twice-daily divided doses, and has a serum level checked for doses above 600 mg/day. Quote that range to its own source. Attributing it to Indian guidance would be a fabrication of provenance, which the reader cannot detect and which no digit check would catch.

IN INDIA

Start lower and climb more slowly than a Western dose table implies. If markedly lower clozapine doses are required in Indian patients, then the same milligram is delivering a higher exposure here than in the populations the published dose ranges were derived in, so the titration should be driven by the response and by dose-related adverse effects rather than by arrival at a printed target dose. If you are titrating without a measured level, that is a reason to move more cautiously: response and tolerability guide the dose, but the exposure itself stays unknown until a valid level is taken, and an unmeasured patient is never a reason to move faster towards the top of a range that was not built on this population.

9.6 What the blood monitoring is actually buying

The obligation that travels with clozapine is haematological, and its size is worth knowing precisely, because candidates habitually overstate it while patients most often refuse the monitoring burden. The Maudsley Prescribing Guidelines put the risk of **agranulocytosis** during clozapine treatment at **1 in 250**, or 0.4 percent, describing that figure as lower than previously thought, and put the risk of death resulting from it, in that same treated population, at 0.05 percent. Around 3.8 percent of patients treated with clozapine develop neutropenia.

1 in 250	0.05%	3.8%
AGRANULOCYTOSIS	FATAL AGRANULOCYTOSIS	NEUTROPENIA

Three figures, three different events, and they are routinely merged into one in an answer. Neutropenia and agranulocytosis are not the same finding and do not carry the same rate; the larger is not a softer statement of the smaller. All three are risks in the treated population, and the third is the one most often misread: fatal agranulocytosis occurs in about 0.05 percent of clozapine-treated patients, and that is not the proportion of those who develop agranulocytosis who die of it. Attaching the figure to the smaller denominator changes what it says. Getting the relation right is worth more than reciting any single value; it is what one follow-up question tests.

The phrase the Maudsley attaches to the agranulocytosis figure, lower than previously thought, is itself examinable: it is a revision of an estimate, not a revision of the monitoring requirement. The requirement itself is set by the guideline bodies, and they do not agree with each other.

9.7 Two blood-monitoring schedules, and where they part

The Indian and United Kingdom schedules open identically and then diverge, which is precisely why one gets substituted for the other without the writer noticing. Indian guidance runs the full blood count and absolute neutrophil count weekly through the first 18 weeks, then monthly for 2 years, then annually, and adds C-reactive protein weekly for a further 4 weeks. United Kingdom practice, as the Maudsley Prescribing Guidelines set it out, runs weekly for the same first 18 weeks, then every 2 weeks to the end of the first year, and usually monthly thereafter.

PHASE OF TREATMENT	INDIAN PSYCHIATRIC SOCIETY SCHIZOPHRENIA GUIDELINE	UNITED KINGDOM PRACTICE, AS THE MAUDSLEY PRESCRIBING GUIDELINES GIVE IT
Before the first dose	Full blood count and absolute neutrophil count, plus C-reactive protein	Full blood count and absolute neutrophil count
Opening stretch	Weekly counts through the first 18 weeks	Weekly counts through the first 18 weeks
Immediately after that stretch	Monthly counts, sustained for the next 2 years	Counts every 2 weeks to the end of the first year
Long term	Annual counts	Usually monthly, with no annual endpoint stated
Inflammatory marker	C-reactive protein at baseline, then weekly for a further 4 weeks	Not part of this schedule

TRAP **Writing one clozapine blood-monitoring schedule as though there were only one.** The two agree for the opening stretch and then separate: Indian guidance steps straight to monthly and holds it for years before annual, while United Kingdom practice inserts a fortnightly phase to the end of the first year and states no annual endpoint. Either can be quoted, provided it is attributed to the body that issues it. Blending the two, or reciting the fortnightly phase in an Indian answer as though it were Indian, is the error, and the answer that names both and marks which is which scores above the answer that picks one and sounds certain.

9.8 When the count falls: lithium, ethnic neutropenia, and the bands this text will not print

A falling count is where the pharmacology of resistance meets the pharmacology of an adverse effect, and it is the point at which a second drug enters for a reason that has nothing to do with symptoms. The Maudsley Prescribing Guidelines list lithium as indicated in patients on clozapine with a low initial white cell count, below 4 x10^9/L, or low initial neutrophils, below 2.5 x10^9/L; and in **leucopenia**, a white cell count below 3 x10^9/L, or a neutrophil count below 1.5 x10^9/L, where the finding is thought to be linked to **benign ethnic neutropenia**.

Two qualifiers on that list carry the weight. The second indication is governed entirely by the benign ethnic neutropenia clause and must never be quoted without it, because it applies where the low count is a constitutional trait rather than a drug effect. And these are values that trigger lithium, not general definitions of leucopenia and neutropenia; the source does not define those terms in this section, and lifting the figures out as definitions misstates what was written.

GOOD TO REMEMBER

Before accepting a handover that says clozapine was stopped for a low count, ask which count, at what value, and whether the patient has a constitutionally low baseline. A low starting count and a constitutional low count are recognised situations with their own management rather than automatically the same event as a fall on treatment, and a patient wrongly closed out of clozapine has lost the option the whole definition of resistance was leading towards.

What this account deliberately does not print is the traffic-light scheme. A green, amber and red banding of the white cell and neutrophil counts, with a defined action attached to each band, is what day-to-day practice actually runs on, and its numbers belong to the product information rather than to either textbook: the marketing authorisation holder's summary of product characteristics and, in India, the approved package insert. Nothing available here carries those bands with their actions attached, and a set of thresholds recalled from memory and printed with consequences beside them is exactly the kind of error that reaches a patient rather than an examiner. Read them off the insert that governs the brand in your own pharmacy, and be suspicious of any revision text that prints them without saying which document they came from.

MNEMONIC

Dose, Duration, Delivery, Dysfunction

The four questions to answer before the word resistant is written. Was the dose pushed to the highest the patient tolerated; was it held for the minimum period the guideline sets for that particular drug; was the drug actually taken and, where it can be, was the exposure measured; and do moderate or severe symptoms still persist alongside psychosocial dysfunction. A trial that fails any of the first three does not belong in the tally, and a picture that fails the fourth is not resistance at all.

Resistance is earned by the trials, not declared by the symptoms: adequate dose, adequate duration, confirmed delivery, measured exposure where monitoring is indicated and available, and persisting symptoms with a functional cost, only then the word.

Drug monitoring, pharmacovigilance and pharmacogenomics

10 · Therapeutic drug monitoring in psychiatry: indications and target ranges (lithium, valproate, clozapine)

Most psychotropics are titrated against the patient. A short list is titrated against a number instead. The marks sit on which drugs belong on that list, when the sample is drawn, and what a printed band licenses; candidates lose them by treating every published range as one range, and by acting on a level that was never a valid measurement.

10.1 Which drugs earn a level, and why

An Indian postgraduate text gives the rationale for **therapeutic drug monitoring** in two limbs: a **narrow therapeutic index**, and confirmation of compliance. Kaplan and Sadock's Comprehensive Textbook of Psychiatry sharpens the first for lithium: it has a narrow therapeutic index, and there is large variability in the oral dose producing a given blood level, arising from differences in renal clearance. For lithium the variability limb is load-bearing, since a narrow margin alone would not compel routine measurement if the dose predicted the concentration.

RULE **A short margin and a dose that predicts the concentration badly are the rationale for measuring lithium routinely.** They are not a pair of conditions every monitored drug must satisfy: confirmation of compliance is a separate indication in the same Indian text, and others follow below. Where a drug has no accepted target range, the Maudsley Prescribing Guidelines state that a plasma level can only indicate adherence or potential toxicity, a two-branch rule that holds for agents no guideline names.

Two further indications sit outside that rationale. The Maudsley guidance on prescribing in physical illness recommends plasma-concentration-guided dosing when narrow-index drugs such as clozapine and lithium are given by enteral feeding tube, and adds, at a distinctly weaker strength, that monitoring may also help in some cases to distinguish non-response from lack of absorption. The underlying principles of half-life and therapeutic index belong to the Psychopharmacology I: principles and core drug classes notes.

10.2 Steady state, and the arithmetic that fixes the first sample

A concentration drawn before the plateau measures where the patient has reached rather than where they will end up, so **steady state** is normally required before a result is read against a target range, and a deliberately early sample has to be labelled as early rather than compared uncritically with a steady-state band. The Maudsley guidelines note that although the common adage is four to five half-lives, three half-lives are sufficient for an approximation of steady state concentrations, and print the ladder that makes this derivable rather than remembered.

75%	87.5%	94%	97%
TWO HALF-LIVES	THREE HALF-LIVES	FOUR HALF-LIVES	FIVE HALF-LIVES

An Indian text gives lithium an average half-life of 18 to 24 hours, from which three half-lives is roughly two and a half to three days and five roughly four to five days. That arithmetic sits behind every published lithium sampling interval. One corollary is examined more often than taught: **loading doses** do not hasten the achievement of steady state.

10.3 The trough, and the conventions that do not agree

A level is a measurement with a time attached. The Maudsley technique for a **trough** sample is to draw immediately before the next dose, never withholding that dose more than one or possibly two hours to obtain it, and the same chapter requires that the time of sampling and of the last dose always be recorded. The first response to a surprising level is therefore to check the timing, not the dose.

> **PITFALL**
>
> Holding the morning dose so the patient can be sampled later in clinic. The result reads low and the reflex is to increase a dose that was never inadequate. The Maudsley limit exists precisely because a longer delay produces a falsely low result.

The conventions diverge, genuinely rather than by error. Kaplan and Sadock time the lithium level as a weekly trough during titration and then every three to six months once stable, whereas the United Kingdom and Indian sources specify hours after the last dose. In twice-daily dosing the two coincide, a sample twelve hours after the last dose falling immediately before the next, and the Indian guideline requires the draw to precede the next dose in twice-daily or thrice-daily regimens. With a single nightly dose that same twelve-hour sample sits mid-interval, a standardised post-dose level and not the predose trough, and that is what a stem supplying a dosing schedule is testing.

SOURCE	WHEN THE FIRST LITHIUM LEVEL IS DRAWN AFTER STARTING OR CHANGING THE DOSE
Indian Psychiatric Society bipolar guideline	At 5 days after dose initiation, sample drawn 11 to 13 hours after the last dose
NICE	At 1 week after starting, 1 week after every dose change, weekly until levels are stable
Maudsley Prescribing Guidelines	At 7 days, then 7 days after each dose change, blood taken 12 hours after the last dose
Kaplan and Sadock	Blood levels determined after 10 days of treatment at a specific dose

Note what the Maudsley entry does not contain. It gives a starting dose and then hands the escalation to the plasma level, repeating that level after every dose change until the desired concentration is reached. No milligram maximum for lithium appears anywhere in the entry,

because the ceiling it works to is a concentration and not a dose. That is the pattern to carry into a lithium question, rather than a top milligram figure the table never issues.

The Indian guideline states its own instruction twice in two different forms: its monitoring footnote gives a window running 11 to 13 hours after the last dose, its narrative a single point at 12 hours. The two are compatible, the point sitting inside the window, but neither is the guideline figure to the exclusion of the other.

10.4 Lithium: bands by phase, by age and by source

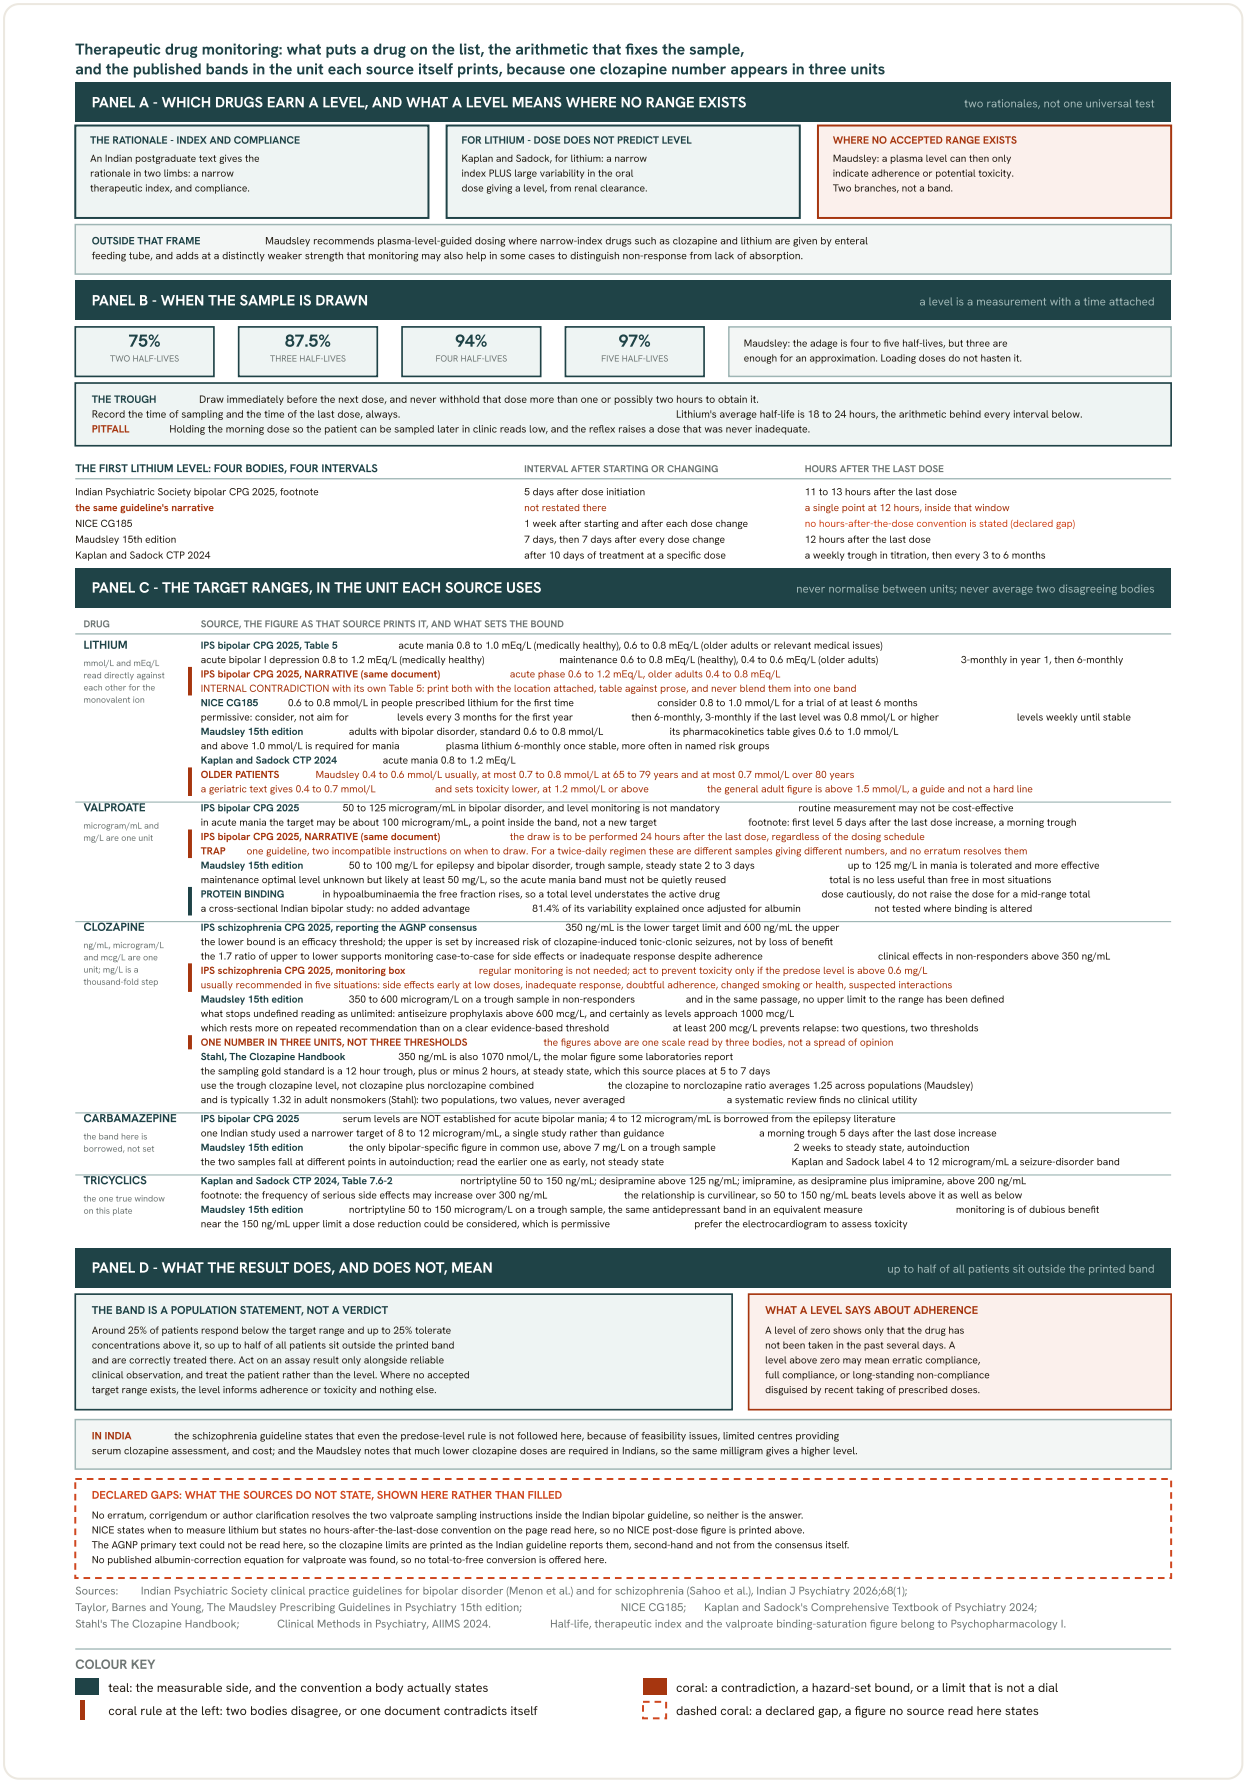


Fig 37.5 Therapeutic drug monitoring: which drugs earn a level, when the sample is drawn, and the bands as each source prints them. The lithium bands differ across four bodies, the Indian bipolar guideline contradicts

itself twice, and the clozapine target is one number printed in three units. (IPS CPG 2025; Maudsley 15th edition; NICE CG185; Kaplan and Sadock, CTP 2024; Stahl, The Clozapine Handbook.)

There is no single lithium target range, and an answer offering one is wrong against three of the four standard sources. The Indian Psychiatric Society bipolar guideline sets acute mania at 0.8 to 1.0 mEq/L in the medically healthy and 0.6 to 0.8 mEq/L in older adults or those with relevant medical issues; acute depression in bipolar type I disorder at 0.8 to 1.2 meq/L in the medically healthy; and maintenance at 0.6 to 0.8 mEq/L in the medically healthy and 0.4 to 0.6 mEq/L in older adults or those with medical comorbidity. That same document then contradicts itself: its narrative on optimising treatment gives a wider acute-phase band of 0.6 to 1.2 mEq/L and an older-adult band of 0.4 to 0.8 mEq/L, which differ from its own table figures. Quoting either alone leaves a candidate wrong against the other.

The other bodies sit near these numbers without landing on them. NICE, under its recommendation on starting lithium, directs aiming for **0.6 to 0.8 mmol/L** in people prescribed lithium for the first time, a lithium-naive rather than phase-specific target, and separately allows clinicians to consider 0.8 to 1.0 mmol/L for a trial period of at least six months after relapse on lithium; note the permissive modal. The Maudsley guidelines give adults with bipolar disorder a standard of 0.6 to 0.8 mmol/L, while their pharmacokinetics table gives 0.6 to 1.0 mmol/L and states that above 1.0 mmol/L is required for mania. Kaplan and Sadock give 0.8 to 1.2 mEq/L for acute mania. For the monovalent lithium ion the millimolar and milliequivalent figures are numerically identical, so these read directly against each other.

In older patients the sources share a floor and diverge on the ceiling. The Maudsley guidelines alone band the target by decade, giving 0.4 to 0.6 mmol/L usually and at most 0.7 to 0.8 mmol/L between sixty-five and seventy-nine years; a geriatric text gives 0.4 to 0.7 mmol/L and defines toxicity lower than the general adult texts, at 1.2 mmol/L or above. The general adult toxicity figure, above 1.5 mmol/L, marks where these ranges stop and is a guide rather than a hard line; toxicity itself belongs to the Psychopharmacology I: principles and core drug classes notes.

Monitoring frequency, once levels are stable, divides on whether the first year is treated differently.

- NICE requires levels every three months for the first year.
- The Indian guideline sets the same three-monthly first year and six-monthly thereafter in healthy individuals, once stable values are seen.
- NICE then drops to six-monthly but keeps three-monthly monitoring for six named higher-risk groups, including anyone whose last level was 0.8 mmol/L or higher.
- The Maudsley guidelines go straight to six-monthly, with more frequent monitoring for interacting drugs, the elderly and renal impairment.

10.5 Valproate: one band, one mania figure, and a contradiction about when to draw

The Indian Psychiatric Society bipolar guideline gives a target of **50 to 125 microgram/mL** in bipolar disorder and states in the same breath that level monitoring is not mandatory, adding that routine measurement may not be necessary or cost-effective in all patients, with poor

therapeutic response or doubts about compliance as examples of when it is indicated. Read the qualifier precisely: not warranted in all patients is not the same as generally unnecessary.

Within that band the guideline notes a linear dose-response relationship in acute mania and states the target may be about 100 microgram/mL, not a separate target. The Maudsley pharmacokinetics table gives a narrower headline range of 50 to 100 mg/L for epilepsy and bipolar disorder on a trough sample, with time to steady state of 2 to 3 days, yet its own comments column concedes that in mania levels up to 125 mg/L are tolerated and more effective than lower concentrations. For maintenance the Maudsley position is an honest blank: optimal levels are unknown but likely at least 50 mg/L, so the acute mania band must not be quietly reused.

TRAP **One guideline, two incompatible instructions on when to draw the valproate sample.** The monitoring footnote permits the first serum level 5 days after the last dose increase, collected in the morning before the first dose of the day. The narrative directs that the draw be performed 24 hours after the last dose, regardless of dosing schedule. For a twice-daily regimen these are different samples giving different numbers, and no erratum resolves them, so a candidate who picks one substitutes their own judgement for the guideline's.

Two further points are frequently mishandled. Valproate is highly protein-bound, so in **hypoalbuminaemia** the **free fraction** rises and a given total level produces a greater effect than expected; the instruction is to dose cautiously, not to raise the dose in response to a mid-range total. Total valproate nonetheless remains the standard analyte, and the sources qualify that rather than universalise it: the Maudsley guidelines state that total concentration is no less useful than free levels in most situations, and an Indian study found free valproate offers no added advantage, with 81.4% of its variability explained once total valproate was adjusted for albumin. The separate binding-saturation figure of 45 to 50 microgram/mL belongs to the Psychopharmacology I: principles and core drug classes notes and is not the bottom of a target band, though the Indian bipolar floor does sit at that saturation point.

10.6 Clozapine: one number in three units

The Indian Psychiatric Society schizophrenia guideline sets the clozapine target at **350 to 600 ng/mL**, reporting these figures second-hand from an international consensus rather than deriving them, and states that non-responders achieve clinical effects when serum levels are above 350 ng/mL. The two ends are set by different things, and that asymmetry is the teachable point: the lower bound is an efficacy threshold, the upper bound, 600 ng/mL, is set by increased risk of clozapine-induced tonic-clonic seizures rather than loss of benefit. With a ratio of 1.7 between its upper and lower target concentrations, clozapine has a narrow therapeutic index window, which the Indian guideline offers as support for monitoring on a case-to-case basis where side effects or inadequate response are suspected despite good adherence, not as a reason to measure every patient.

TRAP **Reading three units as three thresholds.** The Indian guideline's monitoring box triggers action, to prevent toxicity, on a predose level above 0.6 mg/L; the Maudsley pharmacokinetics table gives clozapine 350 to 600 microgram/L on a trough sample; the consensus figures reported by the Indian guideline are 350 ng/mL lower and 600 ng/mL upper. These are one number in three units, not a spread of opinion. Stahl's Clozapine Handbook adds the molar equivalent some laboratories report, giving the response threshold as 350 ng/mL or 1070 nmol/L.

The ceiling is less settled than the schizophrenia band suggests: the Maudsley text states that in non-responders the dose should be adjusted to give plasma levels in the range 350 to 600 microgram/L and, in the same passage, that an upper limit to the target range has not been defined; print it beside the consensus figure, not instead of it. What stops undefined being read as unlimited is the same page's anchor: antiseizure prophylaxis against seizures and myoclonus is prudent when clozapine plasma levels are above 600 mcg/L, and certainly as levels approach 1000 mcg/L, with the caveat that this rests more on repeated recommendation than on a clear evidence-based threshold. Separately, limited data suggest at least 200 mcg/L is required to prevent relapse, so response in schizophrenia and relapse prevention are two questions with two thresholds.

The Indian schizophrenia guideline enumerates five conditions under which clozapine monitoring is usually recommended: early in treatment with prominent side effects at lower doses, inadequate response, doubtful adherence, a change in smoking habits or health, and suspected interactions or side effects. The qualifier is usually recommended, and the box is imported from a United Kingdom practical-considerations paper rather than being guideline-original.

IN INDIA

The same guideline states that regular clozapine monitoring is not needed and is to be conducted only if predose levels are above 0.6 mg/L, especially in smokers, to prevent toxicity, and that in the Indian context even that is not done, because of feasibility issues, limited centres providing serum clozapine assessment, and cost. The Maudsley guidelines note that much lower clozapine doses are required in East Asians, Indians and Bangladeshis, so the same milligram produces a higher concentration here than in the populations the target range came from. Where the assay is unavailable, that argues for a lower starting point and slower titration.

Stahl sets the sampling gold standard as a 12 hour trough, plus or minus 2 hours at steady state, which that source places at 5 to 7 days as a practical assay margin, against the shorter interval the Maudsley table allows for comparable agents. Response correlates best with the trough clozapine level rather than the combined clozapine plus **norclozapine** concentration, which settles which figure the target applies to. The clozapine to norclozapine ratio, not the range-width ratio above, carries two sourced values for two populations and must not be averaged: the Maudsley guidelines give an average of 1.25 across populations with marked individual variation, Stahl a typical 1.32 in adult nonsmokers. A fall in the ratio suggests enzyme induction, a rise inhibition

or a non-trough sample. The counterweight: a systematic review concludes the ratio has no clinical utility.

10.7 Carbamazepine and the tricyclics: a borrowed band and a true window

The Indian guideline is explicit that therapeutic carbamazepine serum levels have not been established for acute bipolar mania, and that clinicians borrow 4 to 12 microgram/mL from the epilepsy literature; printing that as a psychiatric target misstates the guideline. Kaplan and Sadock corroborate the borrowing by labelling the same 4 to 12 microgram/mL band as belonging to seizure disorders. The guideline separately notes one Indian study using a narrower target of 8 to 12 microgram/mL, a single study rather than guidance, while the Maudsley table gives the only bipolar-specific figure in common use, a floor of above 7 mg/L on a trough sample. Timing is where this drug differs: the Maudsley table allows 2 weeks to steady state because of **autoinduction**, whereas the Indian guideline times the sample as a morning trough drawn 5 days after the last dose increase. The two samples fall at different points in that autoinduction, so an early post-increase level need not match a later one and must be read as early rather than as a steady-state result; both are printed.

The tricyclics supply the one true therapeutic window here. The nortriptyline concentration-response relationship is curvilinear, so **50 to 150 ng/mL** is more effective than concentrations above it as well as below it, unlike lithium or clozapine, where the upper bound is set by toxicity. Kaplan and Sadock print nortriptyline at 50 to 150 ng/mL, desipramine above 125 ng/mL and imipramine, measured as desipramine plus imipramine, above 200 ng/mL, with a footnote that the frequency of serious side effects may increase over 300 ng/mL; the Maudsley table gives nortriptyline 50 to 150 microgram/L, the same antidepressant band in an equivalent unit. Only nortriptyline has a ceiling of efficacy, and that is the discriminator. Where the level sits near the upper limit of 150 ng/mL, a dose reduction could be considered, which is permissive rather than mandatory. No Indian guideline addresses tricyclic monitoring, so this window is United Kingdom and United States guidance, not Indian practice.

> **GOOD TO REMEMBER**
>
> The Maudsley guidelines judge tricyclic level monitoring rarely used and of dubious benefit, preferring the electrocardiogram to assess toxicity. Where a tricyclic assay is unobtainable, the tracing is the instrument of choice.

10.8 Reading the result

The interpretive rules carry more marks than the ranges. The Maudsley position is to act on assay results only alongside reliable clinical observation, and to treat the patient rather than the level. The same chapter quantifies it: around 25% of patients respond below the target range and up to 25% tolerate concentrations above it, so up to half of all patients sit outside the printed band and are correctly treated there.

Adherence needs the same discipline. A plasma level of zero indicates only that the drug has not been taken in the past several days. A level above zero may indicate erratic compliance, full

compliance, or long-standing non-compliance disguised by recent taking of prescribed doses, and it is that third possibility a candidate loses by writing that a level confirms adherence. Doubtful adherence is a named indication in both the clozapine box and the valproate footnote, so the limitation must sit beside them.

MNEMONIC

Time, Taken, Target, Tolerance

Four questions before acting on any result: was the sample timed correctly against the last dose, has the drug been taken, does an accepted target range exist for this agent, and is this patient responding and tolerating where they sit. A level failing the first two settles nothing.

A concentration is evidence about exposure, never a verdict on treatment: check the timing, then the patient, then the band.

11 · Pharmacovigilance, adverse drug reaction reporting and signal detection

A licence is granted on a few thousand patients and then tested on a few million. Everything a psychotropic will do that is rare, delayed, or confined to a subgroup is discovered after marketing, by clinicians who notice something and write it down. The marks here are rarely on defining the word. They sit on the discriminations inside it: reaction against event, severity against seriousness, causality grade against causality instrument, and signal against demonstrated harm. Each pair is a place where two words that sound interchangeable carry different technical meanings and send the clinician down a different path.

11.1 The definitions that are examined against each other

An Indian postgraduate pharmacology text gives the wording an examiner expects. **Pharmacovigilance** is defined by the World Health Organization, in its 2002 formulation, as the science and the activities that relate to detecting, assessing, understanding and preventing adverse effects or any other drug related problems. Take those four verbs in order, because they are the architecture of the discipline: nothing can be assessed that was never detected, and nothing can be prevented that was never understood.

Two narrower terms sit inside that definition and are routinely conflated. An **adverse drug reaction**, in that same Indian text's wording, is the technically narrow one: a noxious change suspected to be due to a drug, occurring at doses normally used in man, which requires treatment or a decrease in dose or indicates caution in the future use of the same drug. Three exclusions fall out of that wording, and they are exclusions of this text's definition rather than of the term as every body uses it. On this wording it is not every side effect, because the change must be noxious and must compel a response; it is not poisoning or overdose, because the dose must be one normally used; and it is not a change merely coincident with treatment, because a drug must at least be suspected. The clause requiring that the change compel a response is the textbook's, and it should be attributed rather than carried into a definition credited to anyone else. An **adverse drug event** discards that last requirement, covering any untoward medical occurrence that may

present during treatment with a medicine but does not necessarily have a causal relationship with the treatment.

RULE **Record the event; let the system decide whether it was a reaction.** The event definition is deliberately causality-agnostic and the reaction definition is not. A clinician who waits until they are convinced the drug was responsible has applied the wrong definition at the wrong point in the pipeline, because causality is judged later, by a different body, on the submitted individual case information. Pooled information belongs to signal detection, which is a separate process further downstream.

11.2 Type, severity and seriousness: three classifications, not one

The classification that matters clinically splits reactions by whether they follow from the drug's known pharmacology. **Type B** reactions, the unpredictable or bizarre limb, are based on peculiarities of the patient rather than on the drug's known actions and include allergy and idiosyncrasy; the same text calls them less common, often non-dose related, generally more serious, and requiring withdrawal of the drug. That last property is the operative one at the bedside, because it takes dose reduction off the menu. The receptor logic behind each class's expected adverse-effect signature belongs to the Psychopharmacology I: principles and core drug classes notes; what belongs here is that a reaction outside that signature is a different kind of object.

One sentence in that passage does more work than the rest of the taxonomy: some Type B reactions can be predicted and prevented if their genetic basis is known and a suitable test to characterise the individual's phenotype is performed. That is the argument for pre-prescription testing wherever the genetic basis is known and a validated test exists, and the source confines it to some of these reactions rather than extending it to the Type B limb as a whole. Note too that the sentence asks for a test characterising the phenotype, which genotyping informs without being interchangeable with it. Stated inside a pharmacovigilance chapter rather than a genetics one, it is the hinge to the pharmacogenomics material and to allele-directed testing before the anticonvulsants.

Severity is graded on operational thresholds rather than on impression.

SEVERITY GRADE	WHAT DEFINES IT
Minor	No therapy, antidote or prolongation of hospitalisation is required
Moderate	Requires a change in drug therapy or specific treatment, or prolongs hospital stay by at least one day
Severe	Potentially life-threatening, causes permanent damage, or requires intensive medical treatment
Lethal	Directly or indirectly contributes to the death of the patient

Seriousness is a separate axis with its own outcome list, and it is the one that drives regulatory timelines. A reaction is **serious** if it results in death, a life-threatening condition, hospitalisation

or prolonged hospitalisation, disability, congenital anomaly, an intervention required to prevent permanent impairment or damage, or any other medically important condition. Note what that list does not contain: the reporter's sense of how unwell the patient looked.

TRAP **Writing severity and seriousness as synonyms.** They are independent classifications that happen to share vocabulary. A dystonic reaction settling with one injection took a specific treatment, so it grades moderate on the table above rather than minor, however brief it was; the minor grade needs a change that called for no therapy, no antidote and no added hospital day at all. An admission for observation after a first seizure is serious by the outcome list whatever the reporter thinks of it. Candidates lose the mark by grading one scale and answering on the other.

11.3 Causality: four criteria, and two competing instruments

Causality assessment rests on four named criteria, and every algorithm in use is a way of weighting them. Temporal relationship asks how the time sequence of the event relates to drug administration. Previous knowledge asks whether the drug is known to produce the event in earlier recipients with a certain degree of consistency. **Dechallenge** asks whether the event subsided on stopping the drug, and **rechallenge** asks whether it reappeared when the drug was given again after a gap, a step that is frequently not performed because deliberate rechallenge is unethical or dangerous. The fourth criterion is therefore the strongest evidence available and the least often obtainable.

MNEMONIC

Temporal, Previous knowledge, Dechallenge, Rechallenge

The four questions in the order a case is actually reconstructed: when did it start relative to the drug, is this drug known to do this, did it stop when the drug stopped, and did it return when the drug returned. An answer opening with a scoring instrument has started in the wrong place.

Applied to those criteria, the Indian textbook grades causality into four levels: definite, where causality is proven; probable, where it is not proven but the drug is the likely cause; possible, where the drug as well as other causes could be responsible; and doubtful, where the drug is unlikely to be the cause but cannot be ruled out. The scale used inside the Indian programme is a different one. The World Health Organization and Uppsala Monitoring Centre system carries six categories rather than four: certain, probable or likely, possible, unlikely, conditional or unclassified, and unassessable or unclassifiable. The last two are not degrees of confidence but dispositions for cases that cannot yet be judged, which is why the count differs.

TRAP **Naming one causality instrument as the Indian standard.** Two bodies say two different things and both are examinable. The Indian pharmacology text states that pharmacovigilance centres are expected to assess causality and severity using standard algorithms and rating scales, and offers the Naranjo algorithm for causality assessment and the modified Hartwig scale for severity grading as its examples, not as mandated instruments. The Indian Pharmacopoeia Commission's own programme paperwork states flatly that causality assessment is carried out at the monitoring centres using the World Health Organization and Uppsala Monitoring Centre scale. Print both, attributed. Averaging them, or quietly picking the one you revised, turns a correct fact into a wrong answer.

Two discriminations inside the six-category scale are worth carrying. Certain requires one assessment criterion that probable does not, the fourth: that the event be definitive pharmacologically or phenomenologically. Any other event is automatically excluded and can never qualify for certain, however convincing the time course. Probable and possible separate on a different axis: in possible there may be another equally likely explanation for the event, or there is no information or uncertainty about what happened after stopping.

GOOD TO REMEMBER

In practice few adverse reactions are certain or unlikely; most are somewhere between those extremes, possible or probable. That reflects the instrument rather than the assessor. The source is explicit about the limits: no causality system has been shown to produce a precise and reliable quantitative estimation of the likelihood of the relationship, and the system cannot prove the connection between drug and event, nor change uncertainty into certainty. A causality grade is a structured opinion, worth exactly what its criteria are worth.

11.4 The Naranjo probability scale

The Naranjo adverse drug reaction probability scale is the algorithm most often asked for by name. It is a 10-item questionnaire, and the LiverTox reference that reproduces it gives the arithmetic: total scores range from **-4 to +13**, the reaction being definite at 9 or higher, probable at 5 to 8, possible at 1 to 4, and doubtful at 0 or less. The Indian textbook names the algorithm without printing those bands, so a candidate quoting them should know where they come from.

The internal weighting tells you what the instrument believes. Two items score +2 for a yes rather than the usual one, and they are the temporal item, asking whether the event appeared after the suspected drug was administered, and the rechallenge item, asking whether it reappeared when the drug was readministered; each scores -1 for a negative answer and 0 where the answer is unknown. The weighting is therefore asymmetric rather than doubled in both directions: time and rechallenge gain two points for a yes but cost only the ordinary single point for a no. The instrument rewards a positive time course and a positive rechallenge more than it penalises their absence, which is what makes a negative rechallenge so much weaker evidence than a positive one. Note also that the scale runs to ten items, so it is not the four causality criteria rewritten as arithmetic.

The documentation is candid about the trade. Against the more elaborate organ-specific instrument it is easier to apply but has less sensitivity and specificity in assigning causality in

drug-induced liver injury, and it subtracts points for reappearance of the event on placebo, an item that does not exist outside a trial. A tool carrying an item that cannot be answered outside a trial is being used at the bedside, and that deserves a sentence in any answer that names it.

11.5 The Indian system, and where a report actually goes

The national programme has a datable history that is examined directly. A nationwide pharmacovigilance programme was initiated in **July 2010**, with the All India Institute of Medical Sciences, New Delhi as the first National Coordinating Centre. That role shifted to the Indian Pharmacopoeia Commission at Ghaziabad in April 2011, where it remains, and the commission's coordinating centre was launched as a World Health Organization Collaborating Centre for Pharmacovigilance in Public Health Programmes and Regulatory Services on 30 October 2017. The older textbook description, in which the Central Drugs Standard Control Organization coordinates the programme through peripheral, regional and zonal monitoring centres alongside a national advisory committee, describes the earlier architecture; the current structure is one national centre with a flat network of monitoring centres. Both wordings circulate in question banks.

22	895	4th	3.6%
MONITORING CENTRES AT LAUNCH	CENTRES AT LATEST REPORT	RANK BY REPORT VOLUME	SHARE OF WORLD REPORTS

Date-stamp those figures rather than memorising them loose. The network reached its present size by the end of financial year 2023-24, with 209 centres enrolling in that year alone, and the rank and share above belong to that same index period, measured against 179 member countries of the World Health Organization Programme for International Drug Monitoring. That annual rank is not the country's cumulative standing, which is a separate and lower figure, and an answer quoting one for the other is wrong in either direction.

The pathway is drawable and worth drawing. The monitoring centre reports each individual case safety report to the national coordinating centre through VigiFlow; those reports are assessed at the national centre for quality of data and, if found valid, committed onward to the global drug monitoring centre, the Uppsala Monitoring Centre in Sweden, which is the international collaborating centre. Incomplete or invalid reports go back to the originating centre. VigiBase, the global database at the far end of that pipe, held over 40 million reports from 180+ member countries as of February 2025. A single ward report is one line in that.

11.6 Who reports, on what form, and by when

Reporting duty in India is not one duty, and collapsing it into one is the commonest error in this topic.

REPORTER	INSTRUMENT	WHAT THE SOURCE STATES ABOUT THE DUTY
Patient or consumer	Medicines Side-Effect Reporting Form, available in 10 Indian languages including Kannada, Hindi, Marathi, Tamil and Telugu	The form itself states that reporting is voluntary, has no legal implication, and aims to improve patient welfare
Healthcare professional	Suspected ADR Reporting Form for Healthcare Professionals, Version 1.4, on the commission and regulator websites and in the National Formulary of India 2021	Mandatory fields are patient initials, age at onset of reaction, reaction terms, date of onset, suspected medications and reporter information; incomplete reports are still wanted
Marketing authorisation holder	Expedited individual reports, plus the periodic safety update report under Schedule Y	For a new drug, reports are to be furnished by the applicant every six months for the first two years after approval and annually for the subsequent two years, and the licensing authority may extend the total duration if it is considered necessary in the interest of public health

The asymmetry is the teaching point. All serious adverse events and reactions must be reported by the marketing authorisation holder to the regulatory authority and the national coordinating centre within **15 days** of initial receipt of the information, a hard expedited clock with a named duty-holder. Nothing comparable is stated of the consumer, whose form declares reporting voluntary in terms. Writing that adverse drug reaction reporting is mandatory in India therefore fails on at least one of the three duty-holders, whichever one the examiner had in mind. Beyond the paper form, the programme runs a toll-free helpline on 1800 180 3024, from 9:00 AM to 5:30 PM, Monday to Friday, a reporting mobile application, and an email route to the coordinating centre.

IN INDIA

What to do on the ward when a psychotropic is suspected. Use the Suspected ADR Reporting Form for Healthcare Professionals rather than a free-text note, because its mandatory fields are the minimum the national centre can act on: patient initials, age at onset, the reaction term, the date of onset, the suspected medication and your own details. Send it even when the rest is missing. Do not delay the report while you work towards certainty. The formal grade is assigned at the monitoring centre on the World Health Organization and Uppsala Monitoring Centre scale; your own working judgment about whether the drug did this is for the patient in front of you, not a precondition for sending the form. Where the patient prefers to report independently, the consumer form exists in Kannada as well as nine other Indian languages, which makes direct patient reporting an option to offer rather than a theoretical one.

PITFALL

Holding the report back until you are certain. Two features of the system are built against that instinct: the event definition asks only for an untoward occurrence during treatment and does not require a causal relationship, and the form is built around a small mandatory core with incomplete reports still wanted. The clinician owns the observation and its timing. Formal grading is completed downstream from the individual report; pooled cases are used separately, for signal detection.

11.7 Signal detection: from a database to a hypothesis

A **signal** is defined by the World Health Organization, in the wording the Indian programme itself quotes, as reported information on a possible causal relationship between an adverse event and a drug, the relationship being unknown or incompletely documented previously. The novelty clause is load-bearing, and it turns on novelty rather than on volume: repeated reports of an already-labelled reaction are not new by count alone, though they may still signal a new aspect of the known association. The broader international standard, the CIOMS VIII (2010) definition, widens this in two ways, accepting information suggesting a new potentially causal association or a new aspect of a known association, adverse or beneficial, judged sufficiently likely to justify verificatory action. A new aspect of an old association counts, and so does a beneficial one. What neither definition supplies is causation: a signal does not indicate a direct causal relationship between an adverse event and a medicine, and is essentially a hypothesis.

The statistics are a screening step, not a verdict. The Indian programme performs **disproportionality analysis** on drug-event combinations, based on the contrast between the observed and the expected number of cases, using four key parameters.

- The **Information Component**, which the programme states originated from the Bayesian Confidence Propagation Neural Network and describes as a logarithmic measure of disproportionality, the observed and expected counts being those of one drug and reaction combination against the rest of the report database. The equation itself is deliberately not reproduced. The passage quoted from the programme's report carries the criterion and the credibility interval and does not carry the formula, and the formula as it stands in the report's own text is not a well-formed expression, its numerator and denominator brackets having been lost somewhere between the printed page and the machine-readable text. Reconstructing the brackets would be an inference, and a reader cannot tell a reconstructed equation from a quoted one. Write the Information Component in words unless the primary methodological paper is in front of you. The criterion for a signal is set at **IC025 greater than 0**, IC025 being the lower end point of the 95% credibility interval for each drug and reaction combination, and that criterion is what the source states in its own words. A positive Information Component means the combination is reported more often than the rest of the database predicts, zero means no quantitative dependency, and a negative value means it is reported less often than expected.
- The Proportional Reporting Ratio, with a criterion for a signal of PRR of at least 2.

- Chi-square, with a criterion for a signal of at least 4.
- The number of drug-event combinations, of which at least 3 are required, in the programme's own wording, for the confirmation of a signal. Read the parameter as the count of reports carrying the single drug and reaction pair under test: three is the minimum case count that the disproportionality convention pairs with the ratio and chi-square thresholds, and it is not a demand for three different combinations. Meeting it screens a pair in as a potential signal rather than confirming causality.

No single parameter is decisive: fulfilment of at least three of these four parameters is required before a drug and reaction combination is considered a potential signal. Note also what is being compared. Observed and expected counts are both computed from within the report database, so the comparison sets one drug-event pair against the rest of the reporting, never against the treated population, and no incidence can be read off it.

Human review sits after the arithmetic. The programme's Signal Review Panel comprises scientists and clinical experts affiliated to government and non-government academic institutions and hospitals, and assesses the results of identified computerised signals from individual case safety reports in order to validate and confirm them. The volume surviving to regulatory action is small: in the most recently reported year the programme issued 13 drug safety alerts and recommended 7 prescribing information leaflet changes, including 3 signals, to the Central Drugs Standard Control Organization for appropriate regulatory action.

11.8 Where this touches the psychiatric prescription

The trial phase that finds what matters to psychiatry is the last one. Post-marketing surveillance, Phase IV, provides methods capable of detecting uncommon and idiosyncratic reactions, those appearing only after long-term use, and unsuspected drug interactions. Capable of, not guaranteed to, and not alone in it: the spontaneous reporting stream and the database signal work set out earlier in this section run alongside a formal Phase IV study, as do registries and observational database studies. Its own method is unglamorous: practising physicians are identified and data are collected from them on a structured proforma about the efficacy, acceptability and adverse effects of the drug in the real field situation. Because the study population is patients treated in the normal course, post-authorisation sources can cumulatively cover populations much larger and more representative than the pre-authorisation trials, and that is why they surface what the trials could not, though an individual post-marketing study may be small and an individual licensing programme large. For agents used for years in populations no trial enrolled, this is an important source of the evidence there will be rather than the whole of it.

The prevention list in the same textbook closes the loop with the rest of this chapter. Among its concrete practices is to carry out appropriate laboratory monitoring, with prothrombin time on warfarin and serum drug levels on lithium as its worked examples. A measured concentration is therefore not only a dosing tool but a pharmacovigilance instrument: an adverse reaction prevented rather than reported. Which psychotropics earn a level, when the sample is drawn and

what the printed band licenses are taught where therapeutic drug monitoring is set out, and none of those numbers belong here.

> Report the event, not your verdict on it: causality is graded downstream, and the report that never left the ward is the one the database cannot see.

12 · Pharmacogenomics in psychiatry (CYP2D6, CYP2C19, HLA-B alleles) and clinical application

A genotype changes a prescription only where someone has published what to do with it. The marks in this topic sit on the decision rather than on the molecular biology: which variant is worth knowing before which drug, in which population, and what a result licenses you to change. Two cytochrome enzymes and two HLA loci carry almost all of the examinable material, and candidates lose it in three predictable places. They read a population frequency as an individual phenotype, they carry a recommendation from one drug across to its neighbour on the same enzyme, and they treat a genotype report as a permanent property of the patient when the patient's own prescription has already moved it.

12.1 The two enzymes worth knowing, and what each one clears

KDT ranks **CYP2D6** as the next most important cytochrome isoform after CYP3A4/5, metabolising nearly 20% of drugs, a denominator that is all drugs and not psychotropics, against the nearly 50% biotransformed by the leader. Read that denominator before you quote the figure, because it is the commonest distortion in this topic, and the substrate list the same source gives ranges across tricyclic antidepressants, SSRIs, many neuroleptics, antiarrhythmics, codeine, debrisoquine and metoprolol. An enzyme handling a fifth of the pharmacopoeia is not thereby handling a fifth of your prescriptions.

The Maudsley Prescribing Guidelines print a brief summary of important interactions, and in it the psychiatric substrate cells for the two polymorphic enzymes sit side by side. Read each cell as a selection of important substrates rather than as a complete inventory of what the enzyme clears.

- CYP2D6 metabolises aripiprazole, atomoxetine, brexpiprazole, clozapine, codeine, dextromethorphan, donepezil, duloxetine, haloperidol, phenothiazines, risperidone, tamoxifen, tricyclics, tramadol, trazodone, venlafaxine and vortioxetine.
- CYP2C19 metabolises citalopram, diazepam and moclobemide.

The Maudsley summary table lists three selected CYP2C19 substrates against a much longer CYP2D6 cell, and the short cell is a selection and not the whole of what CYP2C19 clears: the CPIC rows set out below act on escitalopram, which that cell does not name. The short cell is still where the pharmacogenomic decision bites hardest, because citalopram sits on it and citalopram is a psychotropic whose phenotype carries a regulatory dose ceiling. The cytochrome substrate map as a principle, and the interaction matrix built on top of it, belong to the

Psychopharmacology I: principles and core drug classes notes. What follows here is only what a reported phenotype changes.

12.2 From genotype to phenotype: the activity score

A laboratory reports alleles; a prescriber needs a phenotype. For CYP2D6 the bridge is an **activity score**, and CPIC translates that score into a phenotype through a per-score lookup table: 0.0 is a **poor metaboliser**, 0.25 an **intermediate metaboliser**, 1.25 a normal metaboliser and 2.5 an ultrarapid metaboliser. The published table lists a phenotype against each score and does not print the edges between the bands, so the upper edge of the intermediate and of the normal band is not stated here. Two things follow that are worth fixing in memory. Any published poor-metaboliser frequency for this enzyme means the proportion of people sitting at a score of zero and nothing else, so the frequency and the cut-point are one fact rather than two. And the lowest non-zero score is already an intermediate metaboliser under the current standard, not a poor one, which is exactly the boundary a multiple-choice stem will place its distractor on.

The scheme is not shared across genes. For CYP2C19 the phenotype names CPIC uses are poor, intermediate, normal, rapid and ultrarapid metaboliser, so a report that classifies patients into a different set of categories is not reporting the same quantity, however similar the words look.

12.3 How common each phenotype is, and why the group label is half the number

The Maudsley Prescribing Guidelines record that about 20% of Asian people and 3 to 5 percent of people of European descent are CYP2C19 poor metabolisers. CPIC publishes the same quantity by **biogeographic group**, and its figures diverge from those, because the word Asian is undefined in the first source and is split into separate groups in the second.

13.0%	8.2%	4.1%	2.4%
EAST ASIAN	CENTRAL/SOUTH ASIAN	AFRICAN AMERICAN/AFRO-CARIBBEAN	EUROPEAN

Those are CPIC's CYP2C19 poor-metaboliser frequencies, and the South Asian value is roughly three times the European one, which is the single most defensible pharmacogenomic statement available about Indian patients. Within that same Central/South Asian group the rest of the distribution matters as much as its tail: intermediate metabolisers 40.8%, normal metabolisers only 29.6%, and ultrarapid metabolisers 2.9%. A fully normal CYP2C19 phenotype is the minority state in this group, which is not the intuition a table derived from European samples trains.

TRAP **Treating Asian as one population, then averaging two sources that counted different people.** One source quotes a single Asian figure without defining the group; the other splits the same phenotype into East Asian and Central/South Asian and returns two different and lower numbers. No arithmetic reconciles them, because they are not estimates of the same quantity. Print each figure with its population attached, and remember that if the stem says South Asian, the East Asian number is a wrong answer even though it is a correct number.

CYP2D6 runs in the opposite direction, and this is the part most candidates get backwards. CPIC's activity-score-zero frequency is 2.4% in the Central/South Asian group and 6.5% in Europeans, with East Asians lowest at 0.8% and the African American and Afro-Caribbean group at 2.3%. Notice that the Central/South Asian figure for this enzyme is numerically the same as the European figure for the other one in the strip above. Identical digits, different enzyme, different group, opposite clinical meaning: that coincidence is the reason every number in this topic has to travel with both its gene and its population.

TRAP **Assuming slow metabolism travels together.** Poor metabolisers of CYP2C19 are commoner in South Asians than in Europeans; poor metabolisers of CYP2D6 are less common. A stem offering "Indian patients are poor metabolisers" as a single fact is testing whether the candidate knows the statement is enzyme-specific and that its direction reverses between the two enzymes that matter most in psychiatry.

One genuine disagreement should be carried rather than resolved. CPIC's European CYP2D6 poor-metaboliser estimate sits above the figure the Maudsley Prescribing Guidelines give for people of European descent for the same enzyme. Both are published, neither is an error, and the honest answer attributes the number to the body that issued it instead of blending them. What survives both sources is the direction: for CYP2D6 the European rate exceeds the South Asian one.

GOOD TO REMEMBER

These group frequencies are model output, not counts of people. CPIC states that diplotype and phenotype frequencies are estimated from reported allele frequencies using the **Hardy-Weinberg equilibrium** equation, with the margin of error compounding across the alleles that make up a diplotype. Use them as a prior for how likely an unusual phenotype is in the patient in front of you. Do not use them as a substitute for the individual test when the decision genuinely turns on that patient's phenotype.

12.4 What Indian cohorts add, and what a single cohort cannot carry

CPIC's Central/South Asian label is a biogeographic grouping and not a synonym for Indian, so the nearest primary data are individual Indian cohorts. One Indian cohort of healthy volunteers reports a predicted CYP2C19 poor-metaboliser frequency, the phenotype inferred from genotype and not measured, and neither that frequency nor the number of volunteers behind it is printed here: the report itself could not be consulted here, and a prevalence figure is precisely the sort of number a reader carries forward unchecked. The allele frequencies reported for that same cohort are CYP2C19*2 at 0.365, CYP2C19*3 at 0.0033 and CYP2C19*17 at 0.18: one loss-of-function allele that is genuinely common, a second that is vanishingly rare, and a third present at appreciable frequency.

Two disciplines travel with those figures. The first is how to name them. This is one cohort, single-centre and drawn from healthy volunteers, so it is one Indian cohort and never the Indian prevalence, and an answer that promotes it to a national figure has overreached in a way an examiner can mark. The second is how not to combine them. The cohort's phenotypes were

predicted from genotype, while the biogeographic figures were modelled from reported allele frequencies under an equilibrium assumption, and the two were produced by different methods against different phenotype classifications. They can sit in the same discussion; they cannot sit in the same column as though one validated the other.

12.5 What a phenotype actually changes: the gene-drug pair

RULE **The unit of a pharmacogenomic decision is the gene-drug pair, not the gene.** A phenotype means nothing until a body has published an action for the specific drug in front of you, and those actions differ between drugs that share an enzyme, differ between the two ends of one enzyme's range, and are issued at different strengths. CPIC states the principle at its sharpest for the HLA locus: the comment on its carbamazepine recommendation table records that the HLA-B*15:02 test is recommended to guide the use of carbamazepine and oxcarbazepine only. That is the scope of that guideline rather than the whole current scope of the allele, because CPIC issues a separate phenytoin recommendation on the same allele, which is taught with the HLA testing and severe cutaneous reaction material.

PHENOTYPE	DRUG	PUBLISHED ACTION	STRENGTH
CYP2C19 poor metaboliser	citalopram, escitalopram	Consider an antidepressant not predominantly metabolised by CYP2C19; if citalopram or escitalopram remain clinically appropriate, consider a lower starting dose, slower titration and a 50% reduction of the standard maintenance dose	Strong
CYP2D6 poor metaboliser	paroxetine	A 50% reduction in the recommended starting dose, a slower titration schedule, and a 50% lower maintenance dose than in normal metabolisers	Moderate
CYP2D6 ultrarapid metaboliser	paroxetine	Select an alternative drug not predominantly metabolised by CYP2D6	Moderate
HLA-B*15:02 positive, carbamazepine-naive	carbamazepine	Do not use carbamazepine	Strong
HLA-A*31:01 positive, carbamazepine-naive	carbamazepine	Do not use carbamazepine if alternative agents are available	Strong

Read those as five separate instructions rather than as a pattern, because three of the differences between them are examined. The shape of the advice is drug-specific: for the pair metabolised by

CYP2C19 the recommendation leads with choosing a different antidepressant and only then offers a reduced dose, while for paroxetine at the slow end it leads with halving. The strength is drug-specific too, so a classification cannot be carried sideways from one agent to another even when the action is identical. And there is one instruction in this area that comes from outside CPIC's dose adjustments altogether: per the FDA warning quoted by CPIC, citalopram **20 mg/day** is the maximum recommended dose in CYP2C19 poor metabolisers, a ceiling set by the risk of QT prolongation rather than by any limit on antidepressant efficacy. That warning states the ceiling and no starting dose, because a restriction of this kind caps a regimen rather than opening one; where a poor metaboliser is actually being started, the instruction to work from is the CPIC row above, which asks for a lower starting dose and a slower titration without naming a figure for either. The logic underneath all of it is exposure. Less enzyme activity means more parent drug at the same milligram, which is why the published response is a smaller dose or a different drug rather than a different class.

IN INDIA

Because CYP2C19 poor metabolisers are several times commoner here than in the populations most dosing tables were built from, the CYP2C19 row of that table is the one that will alter practice most often. Where citalopram or escitalopram is the chosen antidepressant and a poor-metaboliser phenotype is known, the published instruction is a lower starting dose and a slower titration before any maintenance dose is settled, with the regulatory ceiling applying on top of it. Where the phenotype is unknown and cannot be tested, no phenotype may be read off ancestry: a higher group frequency does not make the patient in front of you a poor metaboliser, and an untested patient is started and titrated on ordinary clinical grounds rather than on a genotype-guided instruction.

12.6 Phenoconversion: when the prescription changes the phenotype

A genotype is fixed; the phenotype it predicts is not. CPIC records that paroxetine-associated **phenoconversion** of CYP2D6 normal metabolisers to intermediate or poor metabolisers, through autoinhibition of the enzyme, may occur, is dose-dependent, and is greater at steady-state concentrations. Read the modal exactly as written: may occur, not does. The clinically useful part is its shape. The drug the enzyme clears is also inhibiting that enzyme, so the phenotype printed on a report taken before treatment is not necessarily the phenotype the patient has once the drug has reached steady state, and the discrepancy grows with the dose.

Autoinhibition is only the tidiest example. The general case is any co-prescribed inhibitor of the same enzyme, which is why a genotype-guided plan has to be read against the current prescription list rather than in isolation. The full inhibitor and inducer matrix across the psychotropic classes is taught with the pharmacokinetic and pharmacodynamic drug interaction material and is not repeated here.

PITFALL

Treating a pharmacogenomic report as a permanent property of the patient. The report describes the genome; the phenotype line on it is a prediction, and that prediction can be moved by the drug being started or by a second drug already in the chart. Before explaining a non-response or an unexpected adverse effect by the genotype, read the prescription list, because a predicted normal metaboliser can behave like a poor one for reasons that have nothing to do with the report and everything to do with what else the patient is taking.

12.7 The HLA alleles: a test used to rule a drug out

The cytochrome tests adjust a dose. The HLA tests remove a drug, and that difference in what the result does is why they are examined together and answered differently. The CPIC HLA-B frequency table gives HLA-B*15:02 a Central/South Asian **allele frequency** of **2.6%**, an East Asian frequency of 4.6% and a European frequency of 0.0000655, that last figure being the table's raw proportion and not a percentage. Two things must be said with those numbers. They are allele frequencies and not carrier frequencies, which is a different quantity that these figures do not give; and the South Asian value, while below the East Asian one, is enormously above the European one, so the familiar teaching that this is an East Asian allele is only half true in an Indian clinic.

The unusual property of the test sits at its negative end. CPIC records that HLA-B*15:02 carries a **100% negative predictive value** for carbamazepine-induced SJS/TEN. That figure is stated flatly in the comment on CPIC's carbamazepine recommendation table, for that reaction with that drug, and it is what makes the allele function as a rule-out rather than a risk score: a negative result greatly reduces the risk of that reaction with that drug rather than abolishing it, and a positive result is a stop rather than a caution. Carbamazepine-associated Stevens-Johnson syndrome and toxic epidermal necrolysis are also reported with other B75 alleles that a single-allele assay does not read, so a negative assay is not a negative patient. It is also why the recommendation is issued with a restriction written into it rather than as a general test of anticonvulsant safety.

HLA-A*31:01 is a second and separate carbamazepine risk allele that does not carry the same consequences as the first, and those differences are precisely what is asked. Its risk spans SJS/TEN, DRESS and MPE, not SJS/TEN alone, so it predicts a wider reaction spectrum including a maculopapular exanthem. Its stopping rule is conditional where the other is absolute: in a carbamazepine-naive patient the instruction is not to use carbamazepine if alternative agents are available. That conditional clause is the discriminator between the two alleles and is worth reproducing in those words rather than paraphrased into a flat prohibition.

For Indian practice the association itself has been quantified. A meta-analysis of Indians from Malaysia and India demonstrated a significant and strong association between HLA-B*15:02 and carbamazepine-induced SJS/TEN, with an odds ratio of 38.54 (95% CI 6.83 to 217.34). The interval teaches as much as the point estimate. The effect is large and the sample it rests on is small, so the finding argues for testing rather than for quoting a precise risk to a patient.

IN INDIA

A low group allele frequency is a poor reason to skip this test, because the frequency describes a population and the reaction happens to an individual. The pharmacogenomic argument runs the other way: the allele is present here at a frequency that is not negligible, the association in Indian patients is strong, and the test's real yield is the negative result it returns to the great majority who are not carriers. Where carbamazepine is the drug being started and the assay is obtainable, that is the moment at which the result changes the prescription. The severe cutaneous reactions themselves, and the wider testing framework reaching phenytoin and oxcarbazepine, are taught with the HLA testing and severe cutaneous reaction material.

12.8 Ancestry, dose requirement and the honest Indian position

Some of what a genotype is invoked to explain is already visible before anyone is genotyped. The Maudsley Prescribing Guidelines record that much lower doses of clozapine are required in East Asians, Indians and Bangladeshis, and that the prevalence of clozapine poor metabolisers is also higher in East Asians. An Indian exploratory study puts a figure on the first half of that: the median clozapine **concentration to dose ratio** in its sample was 2.5 against a published weighted mean of 1.07 for Caucasians, about two and a half times higher.

Both halves of that pairing need discipline. The Indian figure is a concentration observation and not a genotype finding: no variant is named as its cause, and attributing it to a particular cytochrome polymorphism goes beyond what was reported. It is also a single exploratory sample compared against published means rather than against a concurrent control group. What it licenses is the caution that a milligram may not mean here what it meant in the population a dosing table was derived from. The plasma concentration itself, its target range and when the sample is drawn belong with the therapeutic drug monitoring material and are not restated here.

The Indian guideline position is deliberately restrained, and it is the right note to close on. The Indian Psychiatric Society schizophrenia guideline states that **psycho-pharmacogenomic studies**, focusing on cytochrome P450 enzymes, may be helpful in three situations:

- treatment-resistant positive symptoms;
- antipsychotic intolerance;
- frequent relapse despite reported fair medication adherence.

The same guideline then states that the clinical utility of psycho-pharmacogenomics is disputable, because of the lack of larger standardised trials and cost-effectiveness evaluations, and that where such testing is conducted the decisions must always be integrated with other clinical factors. Notice what that position is and is not. It is a judgement about the evidence base and about economic evaluation, and it is stated in a guideline rather than inferred. It is not a statement about which laboratories in the country run these assays, what they cost, or how long they take, and an answer that supplies those details is supplying them from somewhere other than the guideline.

MNEMONIC

Allele, Agent, Ancestry, Action

Four questions before a pharmacogenomic result changes anything: which allele or phenotype has actually been reported, which agent the published recommendation was written for, which population the frequency being quoted was measured or modelled in, and what action a named body has actually published for that pair. A result that fails any one of the four settles nothing.

A pharmacogenomic result is a gene, a drug, a population and a published action: without all four it changes no prescription.

Interactions and prescribing in physical illness

13 · Pharmacokinetic and pharmacodynamic drug-drug interactions of psychotropics

A psychotropic almost never meets a patient alone. The second drug is the commonest second variable here, and three questions are asked of every pair: by what mechanism does one agent reach the other, in which direction does the concentration or the effect move, and how long does the change take to declare itself. Marks are lost by treating every interaction as an enzyme problem, by naming an isoenzyme without committing to a direction, and by assuming a mechanism that is real must also be clinically operative. Half-life, steady state, therapeutic index and the enzymology itself belong to the Psychopharmacology I: principles and core drug classes notes; what follows is the cross-class application, on the enzymatic mechanisms and on the equally examinable mechanisms that have no enzyme in them at all.

13.1 The two axes, and why the second is not the lesser one

Kaplan and Sadock draw the line at the plasma concentration. A **pharmacokinetic** interaction is one drug changing what another drug's plasma concentration turns out to be; a **pharmacodynamic** interaction is one drug changing another's biologic activity. The second definition turns on activity, not on constancy of concentration, so a single pair can sit on both axes at once. Defining that axis as an interaction in which the level does not move imports a condition the standard text does not impose.

RULE **Name the mechanism before naming the drug.** Every interaction here resolves to one of four mechanism families: altered metabolism by a cytochrome or a conjugating enzyme, altered renal elimination, displacement from plasma protein binding, and additive or opposing action at a shared endpoint. The first three move a concentration and are pharmacokinetic; the fourth need not, and is pharmacodynamic.

13.2 Grading an inhibitor and an inducer, and which psychotropics qualify

Strong and moderate are regulatory definitions, not clinical impressions. Both are fixed by what the perturbing drug does to the area under the concentration-time curve of a sensitive **index**

substrate: an inhibitor is graded by how far it raises that exposure, an inducer by how far it lowers it. Induction is graded by loss and inhibition by gain, which keeps the two ladders apart.

5-fold or more	2 to under 5-fold	80 per cent fall or more	50 to under 80 per cent
STRONG INHIBITOR	MODERATE INHIBITOR	STRONG INDUCER	MODERATE INDUCER

The psychotropics at the top of the inhibitor ladder are ordinary drugs. Fluoxetine and paroxetine are the two clinical index strong inhibitors of CYP2D6 named by the FDA, and fluoxetine is a strong inhibitor of CYP2C19 as well, making it a two-enzyme inhibitor rather than a CYP2D6-only one. Fluvoxamine is the index strong inhibitor for both CYP1A2 and CYP2C19, and the same footnote grades it a moderate CYP3A inhibitor and a weak CYP2D6 inhibitor. Three of the commonest selective serotonin reuptake inhibitors therefore carry three different inhibition fingerprints, and knowing the class tells you none of them.

- Carbamazepine is a strong inducer of CYP3A and of CYP2B6 and a weak inducer of CYP2C9.
- The Maudsley Prescribing Guidelines additionally mark carbamazepine and phenobarbital as predominant-pathway inducers of CYP1A2, while the FDA selects no index inducer at all for that isoenzyme; since the FDA table is not exhaustive, that is silence rather than contradiction, and the direction is taught under the Maudsley name.
- No inducer of CYP2D6 is known: the Maudsley inducer column for that isoenzyme reads not known, and it is the only column left empty. Absence of an identified inducer is not a demonstration that the enzyme cannot be induced.

13.3 The substrate side, and why one enzyme reaches several classes

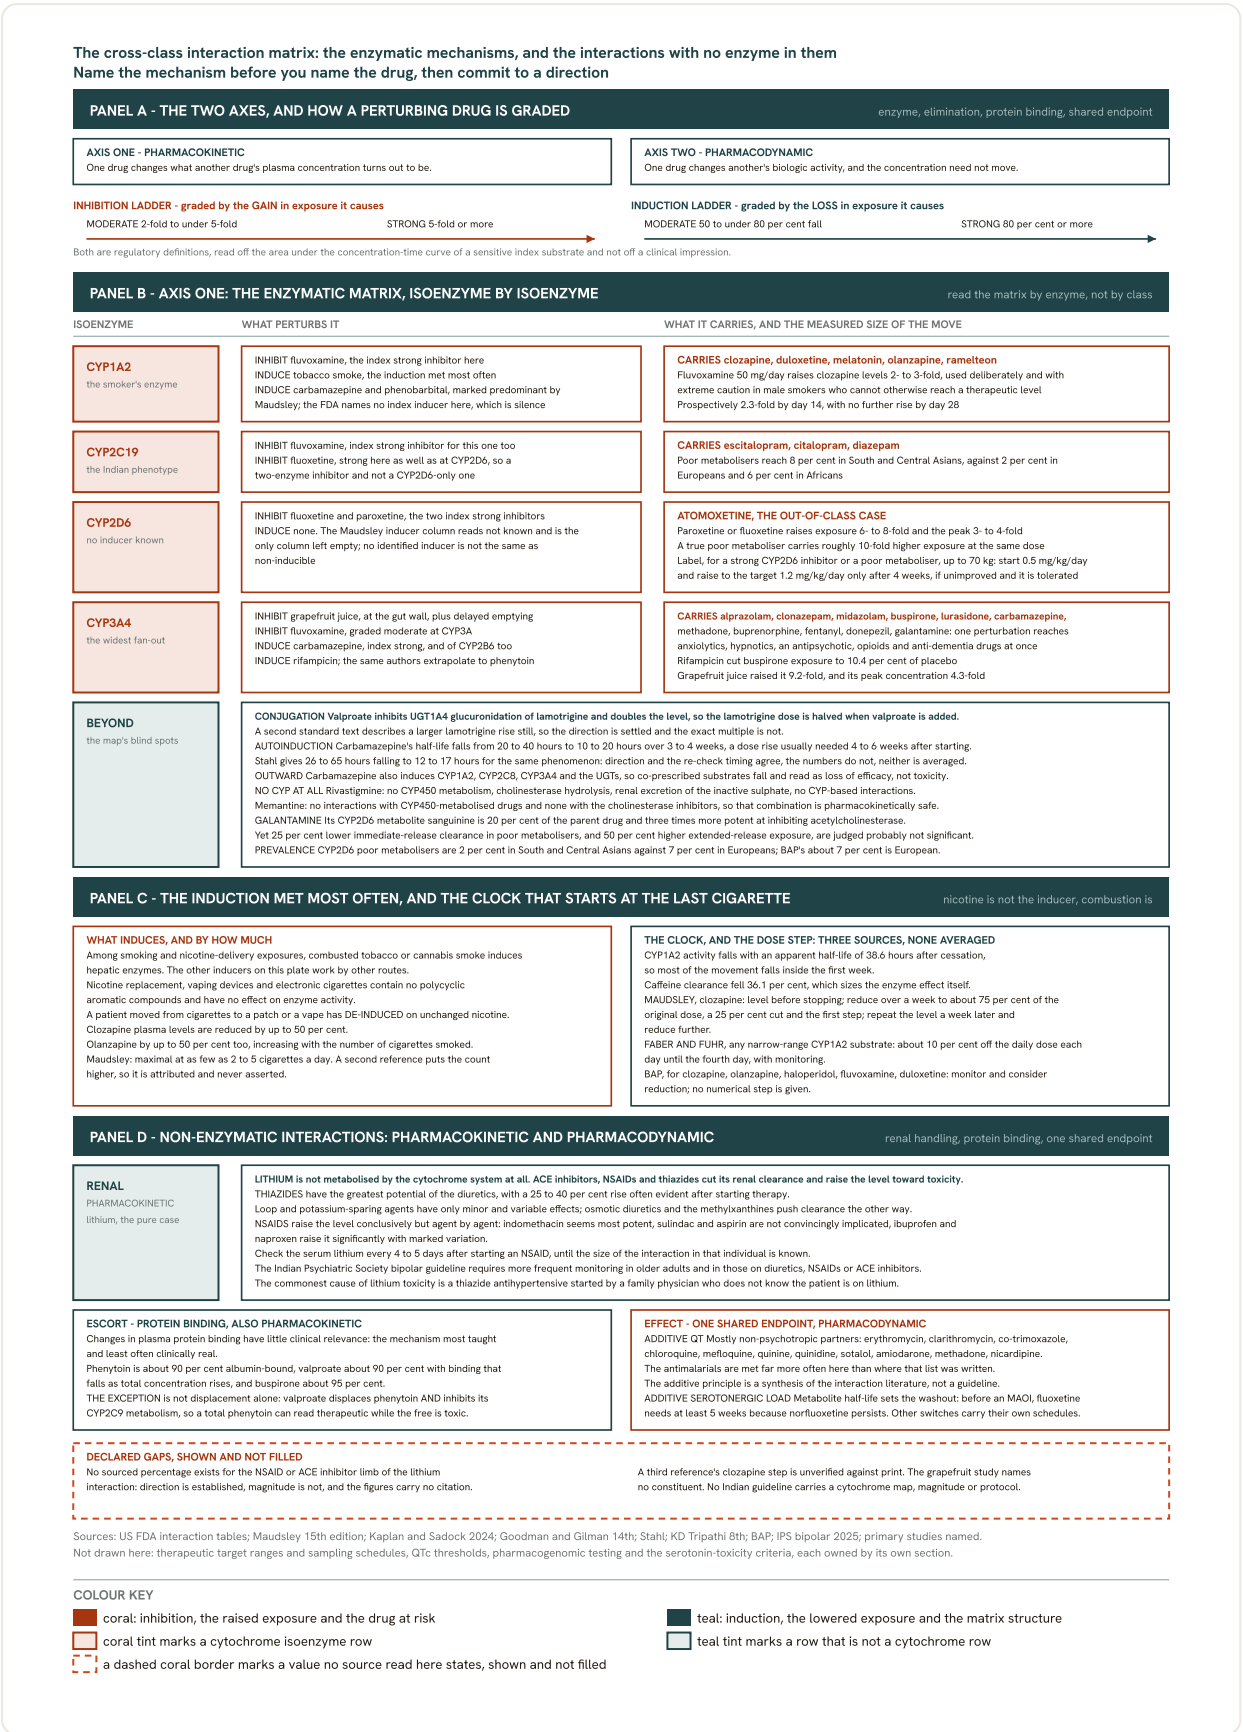


Fig 37.6 The cross-class interaction matrix, on both of its axes. The enzymatic half is read isoenzyme by isoenzyme: what inhibits it, what induces it, what it carries, and the measured size of the move, with conjugation, autoinduction, the agents that have no cytochrome involvement at all, and the population prevalences beneath it. The non-enzymatic half is the one a habit-written answer omits: renal handling with lithium and protein-binding

displacement with its one exception, both pharmacokinetic, and additive action at a shared endpoint, pharmacodynamic.

An inhibitor list is half a prediction. The other half is the substrate map, and a single perturbation fans out widely because the substrates of one isoenzyme are drawn from several therapeutic classes at once. The Maudsley guidelines mark clozapine, duloxetine, melatonin, olanzapine and ramelteon as predominant-pathway CYP1A2 substrates. Their CYP3A4 substrate cell is wider still, carrying alprazolam, clonazepam and midazolam, buspirone, lurasidone, carbamazepine itself, methadone, buprenorphine and fentanyl, and donepezil and galantamine. One CYP3A4 perturbation therefore reaches the anxiolytics, the hypnotics, an antipsychotic, the opioids and the anti-dementia drugs at once, which is the argument for reading the matrix by enzyme rather than by class. The accompanying figure sets out both of its axes together.

13.4 Tobacco smoke, and the clock that starts at the last cigarette

The induction a psychiatrist meets most often is not a drug. Among smoking and nicotine-delivery exposures, the Maudsley position is that combusted smoke induces hepatic enzymes in the manner its own table describes, that this covers cigarettes and cannabis or tobacco joints, and that nicotine replacement, vaping devices and electronic cigarettes do not reproduce that smoke-mediated induction. The footnote attributes that to their containing no **polycyclic aromatic compounds**. Smoking reduces clozapine plasma levels by **up to 50 per cent**, and olanzapine levels by up to 50 per cent as well, that effect increasing with the number of cigarettes smoked. For clozapine the same guideline records that the effect may be maximal at as few as 2 to 5 cigarettes a day; a second standard reference puts the threshold higher, and because the sources disagree the count must be attributed rather than asserted.

TRAP **Treating nicotine as the inducer.** Nicotine replacement, vaping devices and electronic cigarettes contain no polycyclic aromatic compounds and have no effect on enzyme activity, so a patient moved from cigarettes to a patch or a vape has de-induced despite unchanged nicotine intake. The same is true of the smoker admitted to a smoke-free ward. Locating the induction in nicotine predicts no change and misses the rise, which is exactly the stem built around a planned admission or a cessation attempt.

De-induction is quick without being instantaneous. Caffeine phenotyping in heavy smokers put the apparent half-life of the fall in CYP1A2 activity at **38.6 hours**, so most of the movement falls inside the first week; in the same study stopping smoking cut caffeine clearance by 36.1 per cent, which sizes the enzyme effect rather than the effect on any one drug. The participants were few and none were Indian, and that limit travels with the number. Three references address the dose step itself and none agree, so each is attributed and none averaged into the others.

SOURCE, AND THE DRUGS IT SCOPES	THE DOSE STEP ON STOPPING SMOKING	THE MEASUREMENT THAT GOES WITH IT
The Maudsley Prescribing Guidelines, for clozapine	Reduce gradually over a week until around 75 per cent of the original dose is reached, which is a 25 per cent cut and the first step rather than the destination	Take a plasma level before stopping, repeat it a week after stopping, and expect to reduce further
Faber and Fuhr, for any CYP1A2 substrate of narrow therapeutic range	A stepwise daily reduction of approximately 10 per cent until the fourth day after cessation, proposed as a rule of thumb	Accompanied throughout by therapeutic drug monitoring
BAP, for clozapine, olanzapine, haloperidol, fluvoxamine and duloxetine	Monitor and consider dose reduction of the CYP1A2-metabolised drug; no numerical step is specified	Guided by therapeutic drug monitoring
A third standard reference, for clozapine	A different step again, on a different tempo; its figures could not be verified against the printed text and are not reproduced	Not stated here

13.5 Inhibition used on purpose, met by accident, or carried in the genome

Fluvoxamine and clozapine is the case in which the inhibition is the point. The evidence in schizophrenia is that co-prescribing fluvoxamine at a dose as low as 50 mg/day raises clozapine levels 2- to 3-fold, and the combination is used deliberately, with extreme caution, in male smokers who cannot otherwise reach therapeutic levels. The magnitude is confirmed prospectively: mean plasma clozapine rose 2.3-fold at 14 days of combined treatment, with no further elevation at 28 days, so the new plateau arrives inside a fortnight. Used without a measured level, the strategy inverts its own safety logic, because the inhibition is then being relied on blind.

Buspirone demonstrates the same axis in both directions and at full amplitude. Rifampicin cut buspirone exposure to **10.4 per cent of placebo**, effectively abolishing the drug, and the same authors name carbamazepine and phenytoin as potent CYP3A4 inducers with the same consequence, which is their extrapolation rather than a measured comparison. Inhibition runs the other way just as hard: grapefruit juice raised buspirone exposure 9.2-fold and its peak concentration 4.3-fold in healthy volunteers, by inhibiting intestinal CYP3A4-mediated first-pass metabolism and delaying gastric emptying. The primary study attributes the effect to the juice as a whole and identifies no constituent, so naming one is an overreach.

Atomoxetine makes CYP2D6 concrete outside the antidepressants. Adding paroxetine or fluoxetine to a normal metaboliser raises atomoxetine exposure roughly 6- to 8-fold and its peak concentration about 3- to 4-fold, converting the patient pharmacokinetically into a poor metaboliser, and a true **poor metaboliser** carries roughly 10-fold higher exposure than a normal metaboliser at the same nominal dose. The label treats the two conditions as one. In children and adolescents up to 70 kg on a strong CYP2D6 inhibitor such as paroxetine, fluoxetine or

quinidine, or in known poor metabolisers, atomoxetine for attention deficit hyperactivity disorder is started at 0.5 mg/kg/day and increased to the usual target of 1.2 mg/kg/day only if symptoms have not improved after 4 weeks and the starting dose is well tolerated.

Genotype is the standing version of the same inhibition, and the prevalences disagree by source and ancestry. The BAP attention deficit hyperactivity disorder guideline gives CYP2D6 poor metabolisers as about 7 per cent, while the Maudsley guidelines separate the populations and give 2 per cent in South and Central Asians against 7 per cent in Europeans, so the first figure is a European one wearing no label. Allele-level testing and its interpretation are taught in the section of these notes on pharmacogenomics in psychiatry.

IN INDIA

The two polymorphisms point in opposite directions in this population. CYP2C19 poor metabolism reaches **8 per cent** in South and Central Asians, against 2 per cent in Europeans and 6 per cent in Africans, which matters for escitalopram, citalopram and diazepam. Stacking a strong CYP2C19 inhibitor, fluvoxamine or fluoxetine, on one of those agents compounds a common phenotype with a strong inhibitor, so anticipate higher exposure and follow the substrate-specific label or guideline for avoidance, dose reduction, slower titration or monitoring. CYP2D6 runs the other way: poor-metaboliser status is reported less often in South and Central Asians than in Europeans, but genotype and co-prescribed inhibitor both remain clinically relevant and may combine to raise exposure, and the inhibitor is the limb the prescriber can always see.

13.6 Autoinduction, and the drugs with no cytochrome problem at all

Carbamazepine induces its own metabolism. Its half-life falls from 20 to 40 hours down to 10 to 20 hours over 3 to 4 weeks of regular dosing, and the dose usually has to be increased about 4 to 6 weeks after starting to hold the same concentration. Stahl's guide gives different figures for the same phenomenon, an initial half-life of 26 to 65 hours falling to 12 to 17 hours with repeated dosing; direction and the timing of the re-check are agreed across the sources, the numbers are not, and printing either as the figure is an error. **Autoinduction** is also outward-facing: carbamazepine induces CYP1A2, CYP2C8, CYP3A4 and the conjugating UGT enzymes, so co-prescribed substrates of those pathways fall in concentration over the same weeks and present as loss of efficacy rather than as toxicity.

Not every drug has a cytochrome problem. Rivastigmine is the pharmacokinetic outlier of the cholinesterase inhibitors: no CYP450 metabolism, clearance by cholinesterase-mediated hydrolysis with renal excretion of the inactive sulphate metabolite, and therefore no CYP-based drug interactions. Memantine records no interactions with drugs metabolised by CYP450 enzymes and none with the cholinesterase inhibitors, which is what makes the combination pharmacokinetically safe. Galantamine is the counterexample inside the same class: it is metabolised by CYP2D6 and CYP3A4, and its CYP2D6 metabolite **sanguinine** accounts for 20 per cent of the parent drug and is three times more potent than galantamine at producing acetylcholinesterase inhibition. Even there the polymorphism is judged not to matter: Kaplan and Sadock call a 25 per cent decrease in clearance of the immediate-release form in poor

metabolisers, and 50 per cent higher exposure with the extended-release form, probably not significant given the large interpatient variability in galantamine disposition. Hold that beside atomoxetine, where the same enzyme carries the opposite verdict: a polymorphism can be real and still not actionable.

GOOD TO REMEMBER

When a patient with dementia is already established on an enzyme inducer or a strong inhibitor, choosing the cholinesterase inhibitor that is not metabolised by the cytochrome system removes the interaction instead of managing it, and memantine adds nothing to the burden either. That is a decision available at the first prescription, not a monitoring plan bolted on later.

13.7 Non-enzymatic interactions: lithium and the kidney

Lithium is the purest case on this side of the matrix, because it is not metabolised by the cytochrome system at all and its interactions are renal. Kaplan and Sadock state that angiotensin converting enzyme inhibitors, non-steroidal anti-inflammatory drugs and thiazides decrease the renal clearance of lithium and increase the likelihood of severe elevation. The diuretics do not behave alike. Thiazides have the greatest potential of the diuretics to raise lithium, with a **25 to 40 per cent** rise often evident after starting therapy, whereas loop diuretics and potassium-sparing agents have only minor and variable effects, and osmotic diuretics and methylxanthines push clearance the other way and were historically advocated as antidotes for lithium toxicity. No percentage is available for the ACE inhibitor limb: the direction is established and the magnitude is not.

The NSAID limb is agent-dependent in a way the class name hides, and the whole of the account here rests on one paper, named so that a reader can weigh it: Ragheb's review in the Journal of Clinical Psychopharmacology in 1990. Ragheb's review states there is conclusive evidence that these drugs can increase serum lithium levels, diminish renal lithium clearance and possibly induce lithium toxicity, that the effect varies greatly between agents with indomethacin seeming most potent, that sulindac and aspirin are not convincingly implicated, and that ibuprofen and naproxen can raise levels significantly but with marked interindividual variation. No percentage belongs here either: the primary review gives none and the figures in circulation disagree and carry no citation. The review does give a monitoring rule, that a patient on lithium started on one of these drugs should have the serum level checked every 4 to 5 days until the extent of the interaction in that individual is known. The Indian Psychiatric Society bipolar guideline reaches the same place from the other direction, requiring more frequent monitoring in older adults and in those on diuretics, NSAIDs or ACE inhibitors. Target concentrations and the routine sampling schedule are taught in the section of these notes on therapeutic drug monitoring, and lithium toxicity as an entity in the section on the other psychotropic emergencies.

PITFALL

Kaplan and Sadock name the commonest cause of lithium toxicity today as a thiazide antihypertensive started by a family physician who is unaware the patient is under psychiatric care with lithium. It is not created in the psychiatry clinic and cannot be intercepted there, so the warning has to travel with the patient rather than sit in the record.

13.8 Protein binding, conjugation, and the pharmacodynamic axis

Displacement is the mechanism most taught and least often clinically real. The reference position, and the title of the paper that argues it, is that changes in plasma protein binding have little clinical relevance. The figures that make the teaching concrete are unremarkable in themselves: phenytoin is about 90 per cent albumin-bound, valproate about 90 per cent with binding that falls as the total concentration rises, and buspirone about 95 per cent. High protein binding predicts a theoretical interaction, not a clinical one.

TRAP **Reading a therapeutic total phenytoin as a safe phenytoin when valproate is on board.** Valproate displaces phenytoin from albumin and inhibits its CYP2C9 metabolism, so total phenytoin can read therapeutic while free phenytoin is toxic. Displacement alone carries little clinical relevance; displacement combined with inhibition of the same drug's metabolism is the exception that proves the rule, and the Indian standard text makes the same point. Where a free assay is not obtainable, the diagnosis is made from the patient rather than from the report.

Conjugation is the other enzymatic pathway the cytochrome map does not show. Valproate inhibits the UGT1A4 **glucuronidation** of lamotrigine and doubles the lamotrigine level, so the lamotrigine dose is halved when valproate is co-prescribed; a second standard text describes a larger rise still, so the direction is settled and the exact multiple is not.

Two pharmacodynamic stacks carry most of the marks. The first is additive QT prolongation, whose commonest partners in a psychiatric patient are non-psychotropic: erythromycin and clarithromycin, co-trimoxazole, chloroquine, mefloquine and quinine, quinidine, sotalol and amiodarone, methadone and nicardipine. The antimalarials on that list are met far more often here than in the practice it was written for, which changes which partner a psychiatrist encounters. Thresholds and monitoring belong to the section of these notes on QTc prolongation and cardiometabolic monitoring; the additive principle is a synthesis of the interaction literature rather than a guideline position, and is worth stating as such. The second stack is additive serotonergic load, where the arithmetic is set by metabolite half-life rather than by parent half-life: before a monoamine oxidase inhibitor is started after fluoxetine, allow a **washout** of at least 5 weeks because norfluoxetine persists, and an interaction with a serotonergic agent remains possible over that same period after the parent drug has gone. Other serotonergic switches carry their own schedules and do not inherit this interval. Serotonin toxicity itself, with its precipitants, its named criteria and its pharmacological management, is taught in its own section of these notes.

MNEMONIC

Enzyme, Elimination, Escort, Effect

The four families, in the order worth checking them in front of a prescription. Enzyme: is either drug a substrate, inhibitor or inducer of a pathway the other uses. Elimination: does either change the renal handling of the other. Escort: is one displaced from plasma protein binding, with a second mechanism alongside. Effect: do both act on one endpoint, whether or not either concentration moves. A pair on two lines is the dangerous one.

One limitation belongs in any answer built on this material. The enzymatic axis rests on regulatory tables, United Kingdom and United States reference texts and primary pharmacokinetic studies; no Indian guideline carries a cytochrome map, an interaction magnitude or a screening protocol. Say so rather than imply otherwise.

Predict from the mechanism, not from the class: ask which enzyme, which kidney, which binding site and which endpoint the two drugs share, and only then ask which way the number moves.

14 · Psychotropic prescribing in cardiac, hepatic and renal disease

A psychotropic prescription in organ disease is two decisions, not one. The first is ordinary: which drug the disorder needs. The second is what the failing organ does to that drug, or what the drug does to that organ, and the marks are almost entirely for the second. The three organs ask three different questions. The liver principally decides how much of a dose survives to circulate and how fast it is destroyed. The kidney principally decides how much of it leaves. The heart chiefly decides what a given concentration does to conduction, contractility and rhythm, although heart failure can also alter exposure itself through reduced perfusion, hepatic congestion and impaired renal clearance. Candidates lose marks by answering all three with a single list of safe drugs, and by carrying an agent that is the right answer in one organ into the organ where it is the wrong one.

14.1 Reading the organ, not the number

The commonest error at the start is to treat the routine blood test as a measure of the organ's working capacity. The Maudsley Prescribing Guidelines in Psychiatry state plainly that liver function tests are a poor marker of **hepatic metabolising capacity** and not a measure of it: they help evaluate hepatic damage, but tell us little about hepatic function. A patient with grossly deranged transaminases may clear a drug normally, and a patient with a near-normal panel and a shrunken cirrhotic liver may not.

Three organs ask three different questions of one prescription, and a diseased brain asks a fourth: what the failing organ does to the drug, what the drug does to the organ, and which agent is therefore preferred or avoided

PANEL A - THE THREE ORGANS: WHAT CHANGES, WHICH AGENT IS PREFERRED, WHICH IS AVOIDED

the liver test reports injury, not clearance

ORGAN	WHAT CHANGES, AND WHY	PREFERRED OR PERMITTED	AVOIDED OR CONTRAINDICATED
HEPATIC IMPAIRMENT — how much of the dose survives	Most psychotropics are extensively hepatically metabolised, so lower doses may be needed across nearly the whole formulary. Liver function tests index hepatic damage and tell us little about hepatic function: the panel reports injury, not clearance. Monitor LFTs weekly, at least initially; if they deteriorate after a new drug, CONSIDER switching, which is permissive and not directed.	No or minimal hepatic metabolism: sulpiride, amisulpride, lithium, gabapentin. Lithium: no dose reduction if renal function is normal; sample more often when ascites status changes the volume of distribution. IPS tiers: aripiprazole, paliperidone and ziprasidone relatively safer, no adjustment in mild to moderate; olanzapine and clozapine greater risk, risperidone and quetiapine moderate.	Very sedative and very constipating drugs, both for the risk of precipitating hepatic encephalopathy. Hepatotoxic in their own right: MAOIs, chlorpromazine. CONTRAINDICATED: duloxetine, clearance cut even in mild impairment, with fatal hepatic failure reported. Agomelatine, with LFTs at baseline and 3, 6, 12 and 24 weeks and stopping above three times the limit. Valproate in severe or active impairment, greatest liver toxicity risk in the first 3 months.
RENAL IMPAIRMENT — how much of it leaves	The MHRA advises creatinine clearance as the estimate for narrow-therapeutic-index drugs that are mainly renally excreted, lithium among them. Assume at least mild impairment over 65 even when serum creatinine is normal: less muscle mass produces less creatinine. Long-term lithium impairs renal function in about a quarter of patients, and levels above 0.8 mmol/L carry a higher risk of renal toxicity.	IPS: most first-generation antipsychotics are CONSIDERED safe in chronic kidney disease without dose adjustment. Valproate: no routine adjustment, about 2 per cent excreted unchanged. Lamotrigine: under 10 per cent renal excretion of unchanged drug.	Nephrotoxic where renal reserve is limited: lithium, non-steroidal anti-inflammatory drugs. Caution with the extensively renally cleared: sulpiride, amisulpride, lithium. Try to avoid long-acting drugs such as depots. In established renal failure, probably best to avoid the antipsychotics with the GREATEST QTc risk. IPS: phenothiazines may raise hypotension risk, and amisulpride and paliperidone are better avoided in end-stage renal disease.
CARDIAC DISEASE — what the surviving concentration does	In heart failure, SSRIs are the first-line class and sertraline the recommended agent: few drug interactions, less propensity than citalopram to prolong the QTc, and demonstrated safety, though unproven efficacy, in SADHART-CHF. After myocardial infarction SSRIs and mirtazapine have neutral or beneficial mortality associations; the source's do-not-withhold direction is for an SSRI.	Minimal effect on the QT interval and safer at cardiac risk: cariprazine, lurasidone, brexpiprazole, lumateperone. Heart failure: sertraline. Paroxysmal or persistent atrial fibrillation: cariprazine, brexpiprazole or lurasidone may be appropriate, though all antipsychotics appear to raise bleeding risk with direct-acting oral anticoagulants in atrial fibrillation.	Tricyclic antidepressants are generally avoided in cardiac disease: effects on contractility, proarrhythmia from cardiac sodium and potassium channel blockade, and worsening ischaemic heart disease. Higher affinity for cardiac potassium channels, and higher risk of ventricular arrhythmia and sudden cardiac death: haloperidol, olanzapine, quetiapine, risperidone, sulpiride. The recommendation's explicit exclusion: avoid olanzapine.

THE CROSS-ORGAN PIVOT: THE HEPATIC EXCEPTION LIST AND THE RENAL CAUTION LIST ARE THE SAME DRUGS

Sulpiride, amisulpride and lithium are three of the four agents named as undergoing no or minimal hepatic metabolism, which is what makes them attractive in liver disease.
The same three are named as extensively renally cleared, and lithium as nephrotoxic, which is what makes them hazardous in kidney disease.

PANEL B - WHERE THE INSTRUCTION BECOMES A NUMBER

units printed exactly as each source gives them

RENAL: DOSING ACROSS THE GLOMERULAR FILTRATION RATE BANDS, mL/min

Amisulpride is banded by the antipsychotic table and the other three by the mood stabiliser table. The two tables band renal function differently, so each block is headed with its own table's bands.

AGENT	GFR 30 to 60	GFR 10 to 30	GFR under 10
Amisulpride	50 per cent of dose	33 per cent of dose	no recommendation; avoid

AGENT	GFR 10 to 50	GFR under 10
Lithium	AVOID, or 50 to 75 per cent of dose and monitor levels	AVOID IF POSSIBLE; if exceptionally used, 25 to 50 per cent of dose, levels essential
Valproate	no routine adjustment in renal impairment	free (unbound) levels may guide
Lamotrigine	start low and titrate slowly	one source suggests 100 mg every other day

HEPATIC: THE AGENT-LEVEL CEILINGS IN LIVER IMPAIRMENT

Escitalopram, mild to moderate: 5 mg daily for the first 2 weeks, maximum 10 mg daily.
Citalopram (Maudsley): maximum 20 mg/day in hepatic impairment.
IPS: citalopram is not available in India; escitalopram poses no clinically significant QTc risk.
MHRA: escitalopram still causes dose-dependent QT prolongation, so this is not QT safety.
Fluoxetine in compensated cirrhosis: reduce by at least 50 per cent or dose alternate-day.
Venlafaxine and desvenlafaxine: 50 per cent reduction in mild and moderate impairment.
Olanzapine: start 5 mg/day in moderate or severe hepatic impairment; CONSIDER plasma levels, aiming for 20 to 40 mcg/L.
Lurasidone (manufacturer): no adjustment in mild hepatic impairment; start 18.5 mg; maximum 74 mg/day in moderate and 37 mg/day in severe impairment.

ANTIPSYCHOTIC-INDUCED CARDIOMYOPATHY: THE DIRECTION IS THE TABLE'S

45 per cent is a SUGGESTED cut-off, extrapolated from cardiotoxic chemotherapy guidelines.

EF ABOVE 45% →	CONTINUE the effective antipsychotic; 3-monthly symptoms and NT-proBNP
EF BELOW 45% →	SWITCH the antipsychotic.

THE CLOZAPINE DECISIONS IN EPILEPSY

Above a plasma level of 600 mcg/L it is prudent to add antiseizure prophylaxis against seizures and myoclonus, a figure resting on repeated recommendation rather than on a clear evidence base.
After a clozapine-induced seizure: withhold for one day, restart at half the previous dose, and give an antiseizure medication.

PANEL C - THE DISEASED BRAIN: THE SAME THREE QUESTIONS, ASKED OF TWO NEUROLOGICAL ILLNESSES

in both, the first move is subtractive

EPILEPSY — a brain already seizure-prone, and the untreated illness raises the risk too	Depression 13 to 37 per cent, anxiety 20 per cent, psychosis 5 per cent, suicide threefold, and the relationship is bidirectional. Unprovoked seizures about 50 per 100,000 a year, and about 15-fold higher in the PLACEBO arms of psychotropic trials: the untreated illness is itself the seizure risk factor. IPS: ictal and post-ictal psychosis self-remit with supportive treatment and need no prophylactic antipsychotic; inter-ictal and para-ictal often do.	Antidepressants: SSRIs, mirtazapine, duloxetine. Lowest interaction risk with antiseizure drugs: citalopram or escitalopram, then sertraline, with escitalopram first, fewer seizures in overdose. Antipsychotics: amisulpride and sulpiride, aripiprazole, ziprasidone, high-potency first-generation agents, risperidone. With clozapine, lamotrigine is the antiseizure medication of choice. ECT is anticonvulsive and does not worsen epilepsy.	Proconvulsive at therapeutic doses and avoided outright: amoxapine, bupropion, maprotiline, tricyclics. Bupropion's seizure risk is dose related, greatest with the immediate-release form and less on slow release under 300 mg/day. Avoid low-potency first-generation agents, loxapine, depots and zotepine; olanzapine and quetiapine are moderate risk, clozapine higher. Chlorpromazine above 1 g/day: seizures in 9 per cent. Clozapine must not be combined with carbamazepine.
PARKINSON DISEASE — the offending drug is often already on the chart	Psychosis is often characterised by visual hallucinations; on limited data anticholinergics and dopamine agonists carry a higher risk of inducing it than levodopa or COMT inhibitors. Four moves precede the antipsychotic at step 5: exclude organic causes, delirium above all; optimise the environment; do not treat if insight is kept and hallucinations are infrequent and not troubling; and reduce or stop anticholinergics and dopamine agonists.	Only two drugs are recommended for the psychosis: clozapine and pimavanserin, a 5HT2A inverse agonist with no dopamine receptor activity. Low-dose quetiapine is best tolerated and reasonable to try before clozapine. Clozapine: start 6.25 mg, usual 25 to 35 mg/day (Maudsley Table 10.14), monitored as in schizophrenia. Depression: dopamine agonists and SSRIs rank highest for efficacy and acceptability.	Risperidone and typical antipsychotics: avoid COMPLETELY. Pramipexole augmentation MAY be considered, but dopamine agonists raise the risk of impulse control disorders and have rarely been linked to psychosis. ECT: depression and motor symptoms respond well, but delirium risk is high with pre-existing cognitive impairment. Cholinesterase inhibitors: evidence of gains in global functioning, cognition (modest), behaviour and daily activities; motor function may worsen, tremor above all.

DECLARED GAPS: WHAT THE SOURCES BEHIND THIS PLATE DO NOT STATE, SHOWN RATHER THAN FILLED

At EXACTLY 45 per cent ejection fraction the source gives an action above the cut-off and one below it and states no rule at the boundary.
No source here expresses dose adjustment for hepatic impairment as a formal severity score; no Indian availability or price for the QT-safer agents or for pimavanserin; no antidepressant renal band; no millisecond QTc threshold; no NT-proBNP concentration; and no frequency for the lurasidone start, so it prints as 18.5 mg.

Sources: Maudsley Prescribing Guidelines in Psychiatry, 15th edition, Chapters 8 and 10; Maudsley Prescribing Guidelines for Mental Health Conditions in Physical Illness, 1st edition; IPS CPG 2025; MHRA Drug Safety Update on citalopram and escitalopram QTc; Goodman and Gilman, 14th edition. Target ranges, monitoring schedules and the interaction matrix are drawn elsewhere.

teal: the organ's question and the permitted agent · coral: the hazard, the avoid list, the contraindication · coral tint: a decision, not a choice of drug

Fig 37.7 Prescribing in cardiac, hepatic and renal disease, and in epilepsy and Parkinson disease. The liver principally influences how much of a dose is destroyed, the kidney how much leaves and the heart what the surviving concentration does, though each organ works both axes and heart failure may also alter clearance and distribution. Sulpiride, amisulpride and lithium sit on the hepatic exception list and the renal caution list at once. In suspected antipsychotic-induced cardiomyopathy the source table continues above an ejection fraction of 45 per cent and switches below it, the cut-off only suggested and no rule stated at the boundary. In both neurological illnesses the first move is subtractive. (Maudsley Prescribing Guidelines in Psychiatry, 15th edition; Maudsley Prescribing Guidelines for Mental Health Conditions in Physical Illness, 1st edition; Indian Psychiatric Society clinical practice guidelines 2025.)

RULE **The liver test reports injury; it does not report clearance.** Nothing in the standard panel measures how much cytochrome capacity remains, which is why the hepatic dose adjustments below are keyed to a clinical description rather than to a laboratory value: many agent-level adjustments are expressed as mild, moderate or severe impairment, and others as compensated cirrhosis, active disease, ascites, or hepatic impairment generally. These guidelines do not express psychotropic dose adjustment in terms of a formal hepatic severity score at all, so an answer that offers dose bands by severity class is offering something they do not contain.

That matters because **most psychotropics are extensively hepatically metabolised**, so lower doses may be required across nearly the whole formulary. The Maudsley guidelines name the exceptions that undergo no or minimal hepatic metabolism: sulpiride, amisulpride, lithium and gabapentin. Commit that short list, because it is the pivot of the commonest cross-organ trap in this topic.

The renal question is answered differently. The Maudsley guidelines report that the Medicines and Healthcare products Regulatory Agency, a United Kingdom regulator and not an Indian one, advises that **creatinine clearance** be used as an estimate of renal function for direct-acting oral anticoagulants and for narrow-therapeutic-index drugs that are mainly renally excreted, lithium among them. The reasoning is portable even so: the estimate a laboratory prints by default is not necessarily the estimate a narrow-index drug should be dosed against. The underlying kinetics of half-life, clearance and therapeutic index belong to the Psychopharmacology I: principles and core drug classes notes.

14.2 Hepatic impairment: what to avoid, and the watch that follows

Three avoidance rules run across classes in hepatic impairment, and each carries its own reason. Very sedative drugs and very constipating drugs are avoided because of the risk of precipitating **hepatic encephalopathy**, and drugs known to be hepatotoxic in their own right are avoided because the liver is already the injured organ. The Maudsley guidelines give monoamine oxidase inhibitors and chlorpromazine as their examples of the third.

MNEMONIC

Sedating, Constipating, Hepatotoxic

The three properties that disqualify a psychotropic in liver disease, and the first two disqualify it for the same reason: both push an encephalopathy-prone patient towards encephalopathy, one by depressing consciousness directly and one by loading the gut with nitrogen.

The positive instruction is to choose a low-risk drug and monitor liver function tests weekly, at least initially, and it does not direct a switch if they deteriorate: after a new drug is introduced the guideline says only to *consider* switching to another drug. Read that modal exactly. It is permissive, not mandatory, and the reason is the one above: a rising transaminase is a signal about injury, not a verdict on the drug, and the alternative agent may be worse for this patient.

Some sense of how often that signal appears is essential, because an abnormal result is close to routine rather than rare.

One-third	4%	0.5 to 3%
ANY ABNORMAL LFT ON AN ANTIPSYCHOTIC	AT OVER THREE TIMES THE UPPER LIMIT	TRANSAMINASE RISE ON AN ANTIDEPRESSANT

The antipsychotic abnormalities usually appear within 1 to 6 weeks of initiation, and transaminases are most often the affected fraction. The antidepressant elevations are asymptomatic and mild, appear anywhere between several days and 6 months of initiation, and older patients are more vulnerable. Both timings matter, but they change the probability of drug-induced injury rather than settling it: a derangement found at a routine annual test in a patient stable on treatment for years is less likely to be the psychotropic's doing, and one found six weeks in is more likely to be, yet neither latency proves nor excludes the drug. Delayed and chronic presentations occur, so assess the baseline result, the competing causes, the biochemical pattern and the course after withdrawal before attributing the abnormality to the prescription.

14.3 Antidepressants and mood stabilisers in liver disease

Below the class principles sit agent-level instructions, and these are what a stem tests.

AGENT	WHAT HEPATIC IMPAIRMENT DOES TO THE PRESCRIPTION
Escitalopram	Initiate at 5 mg daily for the first 2 weeks in mild to moderate hepatic impairment, maximum **10 mg daily**.
Fluoxetine	Accumulates in compensated cirrhosis; a dose reduction of at least 50 per cent, or alternate-day dosing, is recommended in hepatic impairment, and time to steady state is delayed.
Venlafaxine and desvenlafaxine	Require a 50 per cent dose reduction in mild and moderate hepatic impairment.
Duloxetine	Contraindicated in hepatic impairment; clearance is markedly reduced even in mild impairment and fatal hepatic failure has been reported.
Agomelatine	Contraindicated in hepatic impairment; liver function tests at baseline and at 3, 6, 12 and 24 weeks during initiation and at each dose increase, stopping if transaminases rise more than three times the upper limit of normal.
Valproate	Contraindicated in severe or active hepatic impairment, or where there is a family history of severe impairment, with the greatest risk of liver toxicity falling in the first 3 months of treatment.

The valproate row is the one with the most Indian consequence, because sodium valproate and divalproex are the mood stabilisers most reachable on cost, and alcohol-related liver disease is exactly the population in which a mood stabiliser is often wanted.

TRAP **Two guideline bodies say different things about citalopram, and neither is wrong.** The Maudsley Prescribing Guidelines direct that citalopram's daily dose in hepatic impairment be restricted to a maximum of **20 mg/day**. That entry gives the figure as a restriction on the daily dose and stops; no starting dose accompanies it, so a candidate who recites the ceiling has recited the whole of what the guideline states here. The Indian Psychiatric Society depression guideline states that citalopram can prolong the QTc interval and is not available in India, whereas escitalopram does not pose a clinically significant risk of QTc prolongation. Print that second clause as the guideline's wording, not as a clean bill of health, and correct it as you print it: the United Kingdom medicines regulator's safety advice on the two drugs is that escitalopram also causes dose-dependent QT-interval prolongation. Its effect is generally smaller than citalopram's, but it is not clinically absent, and the usual contraindications and precautions still apply, particularly with cardiac disease, electrolyte disturbance, high exposure or another QT-prolonging drug. A candidate who learns only the British ceiling quotes a limit for an agent the Indian guideline says is not marketed here; one who learns only the Indian statement is caught by a question written from a British text. Attribute each position to the body that issues it and print both.

GOOD TO REMEMBER

Lithium is not hepatically metabolised, so no dose reduction is required in hepatic impairment as long as renal function is normal. The dose does not move, but the sampling discipline does: monitor levels more frequently if ascites status changes, because the **volume of distribution** changes with it. A patient tapped for ascites has not had their dose altered, but they have had their pharmacokinetics altered.

14.4 Antipsychotics in liver disease

Two agents carry usable numbers. The Maudsley guidelines note that although olanzapine is extensively hepatically metabolised, its pharmacokinetics seem to change little in severe hepatic impairment; they nonetheless advise starting at 5 mg/day in moderate or severe hepatic impairment and considering plasma levels to guide dosage, aiming for 20 to 40 mcg/L. For lurasidone, the manufacturer recommends no dose adjustment in mild hepatic impairment, a starting dose of 18.5 mg in moderate or severe hepatic impairment, and a maximum of 74 mg/day in moderate and 37 mg/day in severe hepatic impairment. Set the two rows against each other, because their shapes differ and the difference is the source's rather than an omission here. The olanzapine row fixes where to begin and then hands the rest to a concentration, and it names no milligram maximum at all, which is why the target range and not a dose is the thing to quote from it. The lurasidone row does name maxima, and separate ones for each grade of impairment.

IN INDIA

The Indian Psychiatric Society schizophrenia guideline grades antipsychotic hepatotoxicity in tiers rather than drug by drug, which makes it faster to use at the bedside. Olanzapine and clozapine pose a greater risk and risperidone and quetiapine a moderate risk, while aripiprazole, paliperidone and ziprasidone have relatively safer profiles and need no dosage adjustment in mild to moderate hepatic impairment. That last group is the practical starting point when a patient with liver disease needs an antipsychotic and the prescriber wants one decision rather than a table.

Note the difference in method between the two documents. The British source rates the individual drug and attaches a number to it; the Indian source rates the group and names a safer set. Neither contradicts the other, and an answer that supplies the agent-level ceiling and the group-level preference has covered the question from both directions.

14.5 Renal impairment: which estimate, and in whom

The renal patient is frequently the unrecognised renal patient. The Maudsley guidelines instruct that older adults over 65 be assumed to have at least mild renal impairment even when serum creatinine is normal, because a smaller muscle mass produces less creatinine and therefore a falsely reassuring value. The impairment is present; the marker has simply failed to report it.

> **PITFALL**
>
> Reading a normal serum creatinine in a frail older adult as normal renal function, and leaving a renally cleared psychotropic at the dose that suited the patient a decade ago. The prescription has not changed and the number on the report has not changed, so nothing prompts review, and the concentration climbs anyway. The same failure occurs in reverse when renal function falls acutely and the drug chart is not revisited.

Three general principles then organise the choice of agent:

- Avoid drugs that are **nephrotoxic**, lithium and non-steroidal anti-inflammatory drugs among them, where renal reserve is limited.
- Be cautious with drugs that are extensively renally cleared, such as sulpiride, amisulpride and lithium.
- Try to avoid long-acting drugs such as depot preparations, which cannot be withdrawn once administered.

A fourth principle is to try to avoid drugs known to prolong the QTc interval, and it is narrowed usefully for the sickest patients: in established renal failure, where electrolyte changes are common, it is probably best to avoid the antipsychotics with the greatest risk of QTc prolongation, not every drug that prolongs it at all. Read that as written: a prohibition on every agent with any QTc effect would exclude almost the entire class.

TRAP **The hepatic exception list and the renal caution list are the same drugs.** Sulpiride and amisulpride are named as escaping hepatic metabolism, which makes them attractive in liver disease, and they are named again as extensively renally cleared, which makes them hazardous in kidney disease. Lithium sits on both lists for the same reason. A candidate who memorises one list and reaches for it in the other organ has been set up by the pharmacology itself: a drug avoids one route of elimination precisely by depending on the other.

14.6 Dosing by glomerular filtration rate

Where the renal instruction becomes quantitative, it is banded rather than continuous.

AGENT	RENAL HANDLING	DOSING ACROSS THE BANDS
Amisulpride	Extensively renally cleared	In renal impairment, 50 per cent of the dose at a GFR of 30 to 60 mL/min and 33 per cent at 10 to 30 mL/min; no recommendation exists below 10 mL/min, so it is best avoided in established renal failure.
Lithium	Nephrotoxic; 95 per cent excreted unchanged in urine, and contraindicated by the manufacturer in severe renal impairment	At a GFR of 10 to 50 mL/min, avoid, or reduce to 50 to 75 per cent of the normal dose and monitor serum levels; below 10 mL/min, avoid if possible, and only if it is used exceptionally reduce to 25 to 50 per cent, with level monitoring essential.
Valproate	Approximately 2 per cent excreted unchanged	Dose adjustment is usually not required in renal impairment; only in severe impairment, with a GFR under 10 mL/min, may it be necessary to dose against **free (unbound) valproate levels**.
Lamotrigine	Under 10 per cent renal excretion of unchanged drug	At a GFR of 10 to 50 mL/min start low and titrate slowly in renal impairment; one source suggests 100 mg every other day where the GFR is under 10 mL/min. No lamotrigine starting dose is named at either level of function, and no maximum at the lower one.

Lithium also generates the renal problem it is dosed around. Long-term treatment may result in impaired renal function in about a quarter of patients, the phenomenon usually called **creatinine creep**, and levels above **0.8 mmol/L** are associated with a higher risk of lithium-related renal toxicity. That threshold is a prescribing constraint rather than a laboratory curiosity, because it sits inside the range some phases of illness are treated in; the published target concentrations themselves are set out elsewhere in these notes.

IN INDIA

The Indian Psychiatric Society schizophrenia guideline takes a more permissive line on the older agents than the British drug-by-drug tables do, and this matters where first-generation antipsychotics carry the public-sector caseload. Most first-generation antipsychotics are considered safe in chronic kidney disease without dose adjustment, the exception being the phenothiazines, which may increase hypotension risk. Among the second-generation agents, amisulpride and paliperidone are excreted unchanged in urine, need dose consideration, and are better avoided in end-stage renal disease. Keep the guideline's own hedges in the answer: *most* first-generation agents, and *considered* safe.

14.7 Cardiac disease: the antidepressant decision

The cardiac question is predominantly pharmacodynamic, where hepatic and renal dose adjustment is predominantly pharmacokinetic, though each of the three organs involves both axes, and the class answer is settled for one setting. In heart failure, selective serotonin reuptake inhibitors are the first-line antidepressant class and the recommended agent is sertraline: it has few drug interactions, less propensity than citalopram to prolong the QTc, and demonstrated safety, though unproven efficacy, in SADHART-CHF. Safety without proven efficacy is the honest reading and should be said aloud in an answer rather than upgraded. The recommendation behind this sentence is written for heart failure and reaches no further, so whether the same class ranking holds across cardiac disease at large is a question these notes leave open rather than answer.

Tricyclic antidepressants are generally avoided in patients with cardiac disease for three separate reasons, and separating them is what distinguishes a good answer from a list:

- their effects on cardiac contractility;
- their proarrhythmic effects, arising from blockade of cardiac sodium and potassium channels;
- their potential to worsen ischaemic heart disease.

This is the single most consequential exclusion in an Indian outpatient department, where the tricyclics remain heavily used on grounds of cost. The pharmacology that produces those three effects, and the class's membrane-stabilising action, belongs to the Psychopharmacology I: principles and core drug classes notes; what belongs here is that the cardiac history, not the depressive presentation, is what removes the class from consideration.

The converse error is equally examinable. After myocardial infarction, selective serotonin reuptake inhibitors and mirtazapine have either a neutral or a beneficial effect on mortality, and the instruction that follows from it is written for the first of those two: treatment of depression with a selective serotonin reuptake inhibitor should not be withheld post infarction. The mortality finding covers both agents; the directive not to withhold does not, so do not extend it to mirtazapine in an answer. Depression after an infarct is not a reason to defer antidepressant treatment until the heart has settled.

14.8 Cardiac disease: the antipsychotic decision

Antipsychotic choice in cardiac disease turns on two distinct liabilities: the effect on the QT interval, and the effect on the myocardium itself. On the first, cariprazine, lurasidone, brexpiprazole and lumateperone are considered safer choices in patients at risk of cardiac events, as they appear to exert minimal effects on the QT interval, and the recommendation carries an explicit exclusion: avoid olanzapine. Against them sit haloperidol, olanzapine, quetiapine, risperidone and sulpiride, which have higher affinity than other antipsychotics for cardiac potassium channels and are associated with a higher risk of ventricular arrhythmia and **sudden cardiac death**. Three of the preferred agents recur in a second setting: in paroxysmal or persistent atrial fibrillation, cariprazine, brexpiprazole or lurasidone may be appropriate choices,

and all antipsychotics appear to increase the risk of bleeding when combined with direct-acting oral anticoagulants in atrial fibrillation.

The second liability is **antipsychotic-induced cardiomyopathy**, and here the decision is not which drug but when to stop it. A left ventricular ejection fraction of **45 per cent** has been suggested as the cut-off, and the source's own table fixes the direction: while the ejection fraction stays above that figure the effective antipsychotic is continued, with 3-monthly monitoring of heart failure symptoms and NT-proBNP, and once it falls below that figure the antipsychotic is switched. Read the direction off the table rather than off the surrounding prose, because a falling ejection fraction is the signal to come off the drug and never the licence to stay on it. The word suggested is load-bearing, since the figure is extrapolated from guidelines for monitoring cardiotoxic chemotherapies and the exact level depends on the clinical scenario, so an answer that presents it as an established cut-off overstates it. The table states an action above the cut-off and an action below it and gives no rule at the cut-off itself, so the boundary case is left open here rather than answered.

> The liver principally influences how much drug is destroyed, the kidney how much leaves, and the heart what the surviving concentration does, with heart failure able to alter clearance and distribution as well; the agent that escapes one of the first two organs almost always depends on the other.

15 · Psychotropic prescribing in epilepsy, Parkinson disease and other neurological illness

The second variable here is a brain that is already diseased. Prescribing in epilepsy and in Parkinson disease inverts the ordinary logic of psychotropic choice. The agent that best fits the psychiatric syndrome may destabilise the neurological one; the drugs already on the chart for the neurological illness may be producing the psychiatric syndrome; and the option that looks safest, withholding psychiatric treatment altogether, carries a risk of its own. The marks sit on which agent, in what order, and on what is done before anything is written at all.

15.1 Epilepsy and psychiatric illness travel in both directions

People with epilepsy carry an elevated prevalence of several psychiatric disorders, and the Maudsley Prescribing Guidelines are explicit that the relationship with mental illness is **bidirectional** rather than a simple consequence of living with seizures.

13 to 37%	20%	5%	3x
DEPRESSION	ANXIETY	PSYCHOSIS	SUICIDE

Bidirectionality is what the examiner is testing, because it changes the direction of the clinical question. If the psychiatric disorder were only a reaction to the burden of a chronic illness, treating it would be a humane addition to seizure control. If the two illnesses share substrate and each raises the risk of the other, treating the psychiatric disorder becomes part of managing the epilepsy.

A second observation sharpens it. The incidence of unprovoked seizures in the general population is about 50 per 100,000 persons a year, while the incidence in the placebo arms of randomised controlled trials of antidepressants and antipsychotics is approximately 15-fold higher than that yearly figure. Nobody in a placebo arm has received an active drug, so the excess cannot be pharmacological: depression and psychosis are themselves seizure risk factors.

TRAP **Concluding that the safe course in epilepsy is to withhold the antidepressant.** Candidates reason correctly that several psychotropics lower the seizure threshold, then draw the wrong conclusion from it. The placebo-arm figure settles the argument the other way: untreated depression and psychosis carry a seizure incidence far above the population background, so leaving the psychiatric illness untreated is not the neutral option it appears to be. The decision is which agent, not whether.

15.2 Treating depression in epilepsy

The Maudsley Prescribing Guidelines sort antidepressants for people with epilepsy into low-risk good choices and a higher-risk group that should be avoided because its members are **proconvulsive** at therapeutic doses. The good choices are SSRIs, mirtazapine and duloxetine; the group to avoid is amoxapine, bupropion, maprotiline and TCAs. Read the qualifier on the second group, because it carries the whole instruction: these agents provoke seizures at ordinary treatment doses, not only in overdose, which is why they are avoided outright rather than merely capped.

MNEMONIC

In: SSRIs, Mirtazapine, Duloxetine. Out: Amoxapine, Bupropion, Maprotiline, Tricyclics

Three in, four out, and the four out share a single reason. Nothing else needs to be memorised about the antidepressant list for epilepsy; everything else is a decision about which of the three to use.

Within the SSRIs the discriminator is interaction, not seizure risk. Those with the lowest risk of interactions with antiseizure medications are generally preferred, which gives citalopram or escitalopram, followed by sertraline, and escitalopram is preferred over citalopram in people with epilepsy because of the lower risk of seizures in overdose. Overdose risk is a legitimate selection criterion in a population whose suicide rate is a multiple of the general population's. The kinetic principle that generates the interaction ranking belongs to the Psychopharmacology I: principles and core drug classes notes.

IN INDIA

Citalopram is not marketed here, so the preference order collapses without loss to escitalopram, then sertraline. For antipsychotics, the Indian Psychiatric Society schizophrenia guideline directs that where one is needed an agent of lower seizure propensity is chosen, and it groups quetiapine and olanzapine with clozapine, chlorpromazine, loxapine and zotepine as agents to avoid or use with caution. That is one tier stricter than the Maudsley position, which calls quetiapine and olanzapine moderate risk. Followed together, the two documents leave a practical Indian starting set of amisulpride or sulpiride, aripiprazole, ziprasidone, risperidone and the high-potency first-generation agents.

Bupropion is the instructive member of the avoid group, because its risk belongs as much to the preparation as to the molecule: the seizure risk is dose related and greatest with immediate-release formulations, and it is less with slow-release formulations at doses **under 300 mg/day**, which is where the seizure evidence puts the lower risk of the two preparations. That does not lift it out of the higher-risk group for a person who already has epilepsy. It explains why one drug can be rated differently in different populations, and it is the reason a formulation must be specified whenever this agent is discussed.

Where the depression is severe and the epilepsy unstable, the modality question arrives early rather than late. Electroconvulsive therapy has anticonvulsive properties, is worth considering for depression in patients with unstable epilepsy, and does not appear to cause or worsen epilepsy. That is the opposite of the intuition most candidates carry into a viva, and it is worth a sentence in any answer on this ground. The procedural technique is set out in the ECT and neurostimulation notes.

15.3 Antipsychotics in epilepsy: the tiers

The same source ranks antipsychotics on seizure propensity in four tiers. The low-risk good choices are amisulpride and sulpiride, aripiprazole, ziprasidone, high-potency first-generation antipsychotics and risperidone; olanzapine and quetiapine are moderate risk with care required; clozapine is higher risk; and the low-potency first-generation agents, loxapine, depots and zotepine should be avoided. The depots earn their place in the avoid group for a kinetic reason that generalises beyond this table: an agent that cannot be withdrawn quickly is a poor choice when the adverse effect you fear is a seizure.

One figure anchors the avoid tier. In people with epilepsy the low-potency first-generation agents are best avoided, and the incidence of seizures with chlorpromazine is **9 per cent at doses above 1 g/day**, an incidence the evidence attaches to that dose band rather than a target to prescribe into. This is not a historical curiosity in Indian practice, where high-dose chlorpromazine is still encountered on general wards, so the number is worth carrying.

15.4 Clozapine in epilepsy: three decisions

Clozapine is the most proconvulsive antipsychotic, and the same source records that it has still been used successfully in people with epilepsy who are stable on antiseizure medication. The risk is therefore managed rather than avoided, and three decisions follow, examined together.

The first is co-prescription. Clozapine should not be used with carbamazepine, because of the risk of blood dyscrasias and reduced clozapine levels, and lamotrigine is the antiseizure medication of choice alongside it. Two independent mechanisms converge on one instruction there, one haematological and one pharmacokinetic, which is why the prohibition is firmer than an ordinary interaction warning.

The second is prophylaxis. The Maudsley advises that it is prudent to use an antiseizure agent as prophylaxis against seizures and **myoclonus** when clozapine plasma levels are above **600 mcg/L**, and states in the same breath that this figure is not a clear evidence-based threshold but a level resting on repeated recommendation. Print the caveat with the figure. A reader who knows the source knows the number is a convention, and an answer that treats it as a hard biological boundary reads as second-hand.

The third is what happens after a seizure. Clozapine is withheld for one day, restarted at half the previous dose, and an antiseizure medication is given. Notice that the drug is not abandoned: the response is a brief interruption, a halved reintroduction and added cover, which is a dosing decision rather than a discontinuation. The immediate care of the patient who is fitting belongs to the Perinatal/women's mental health and emergency psychiatry notes.

> **PITFALL**
>
> Reaching for carbamazepine as the antiseizure cover. It is the reflex choice on a general ward and it is the single agent named against this combination, on two counts at once: a shared risk of blood dyscrasias, and a fall in the clozapine level that can be misread later as non-adherence. Lamotrigine is the named alternative, and the substitution should be made before the seizure, not after it.

15.5 The psychoses of epilepsy: which ones need an antipsychotic

Not every psychosis in a person with epilepsy is a psychotic illness, and the Indian Psychiatric Society schizophrenia guideline makes that discrimination the first step. Ictal and **post-ictal psychosis** are to be ruled out: they are self-remitting with supportive treatment, and preventing seizures would suffice without the need for prophylactic antipsychotic treatment. **Inter-ictal psychosis** and para-ictal psychosis would have a course similar to schizophrenia and often need antipsychotic treatment.

Read the qualifier on the first limb, because it is load-bearing: self-remitting with supportive treatment is not the same as self-remitting if left alone. The two standard sources do not say this identically. Where the Indian guideline holds that seizure control suffices without prophylactic antipsychotic treatment, the Maudsley adds a short-term symptomatic option, a benzodiazepine or an antipsychotic, tapered off carefully once symptoms resolve. They are not reconciled by picking one. The Indian wording is the citable position for an Indian answer; the Maudsley line is the comparator a strong answer names beside it.

TRAP **Treating every psychosis in epilepsy as schizophrenia and starting maintenance antipsychotic treatment.** The stem usually supplies the discriminator and candidates read past it: a psychosis tied to the seizures, which settles as the seizures are controlled, is ictal or post-ictal, and the guideline position is that preventing seizures suffices without prophylactic antipsychotic treatment. The error is not which drug was chosen. It is treating the episode as one that needs maintenance: short-term symptomatic treatment with a benzodiazepine or an antipsychotic may still be required and is tapered off carefully once symptoms resolve, and what the guideline position excludes is the prophylactic course that is started and then continued.

15.6 Psychosis in Parkinson disease: what precedes the prescription

In Parkinson disease the offending drug is usually already on the chart. Psychosis in Parkinson disease is often characterised by **visual hallucinations**, and the drug ranking that follows rests on limited data and is not a superlative about any single agent: anticholinergics and dopamine agonists seem to be associated with a higher risk of inducing psychosis than levodopa or catechol-O-methyltransferase inhibitors. It ranks one pair of drug groups against another, which makes the chart the first thing to read.

The Maudsley recommendations for the treatment of psychosis in Parkinson disease therefore place four moves ahead of any antipsychotic.

- Exclude organic causes, delirium in particular.
- Optimise the environment, to maximise orientation and minimise problems arising from poor caregiver-patient interactions.
- If the patient has insight and hallucinations are infrequent and not troubling, do not treat.
- Consider reducing or stopping anticholinergics and dopamine agonists.

Only after those does the ladder reach the step that considers an atypical antipsychotic. Three of the four are subtractive or observational, and the third is the one candidates omit: a treatment decision can be a decision not to treat, and the source states that as an instruction rather than as a permission.

RULE **In both illnesses the first move is subtractive.** In epilepsy the question is which psychotropic is provoking, or would provoke, the seizure. In Parkinson disease it is which dopaminergic or anticholinergic drug is producing the psychosis. Adding a psychotropic to a neurological illness is the second decision and never the first, and an answer that opens with the antipsychotic has skipped the half that carries the marks.

15.7 Which antipsychotic, when one is needed

Low-dose quetiapine is the best tolerated antipsychotic in Parkinson disease psychosis and it is probably reasonable to try it before clozapine, while risperidone and typical antipsychotics should be avoided completely. The word to notice is completely: this is not a caution about dose or duration but an exclusion of two whole groups, and the reason is mechanistic, since dopamine receptor blockade is the therapeutic action in schizophrenia and the pathogenic one here.

Clozapine for psychosis in Parkinson disease is started at 6.25 mg with a usual dose of **25 to 35 mg/day** in the Maudsley table of recommendations for treating psychosis in Parkinson disease. Name that table whenever the figures are printed, because the same book's guide to medication doses in older adults gives a different usual range for the same indication and adds a stated daily maximum the psychosis table does not carry. The two cells are not reconcilable, and quoting one without naming its table quotes the book against itself. Say the consequence plainly when you use the psychosis table: it gives a start and a usual range for this indication and states no ceiling above them, so the top of that range is where the usual dose ends and not where the drug is capped.

GOOD TO REMEMBER

Clozapine in Parkinson disease psychosis is monitored exactly as clozapine in schizophrenia is monitored, and older people are more prone to develop serious blood dyscrasia. A dose in the tens of milligrams does not buy a lighter haematological schedule. That, rather than the price or the supply of the drug, is what limits its use here, and it is the question to settle before the first prescription rather than after it.

Two drugs, and only two, are recommended for Parkinson disease psychosis: clozapine and pimavanserin, a 5HT2A receptor inverse agonist with no dopamine receptor activity. A network meta-analysis found only these two efficacious while having minimal effect on motor function. The mechanism explains the pairing, since an agent that suppresses psychosis without occupying the dopamine receptor has a low propensity for direct dopamine-blockade-related motor worsening. Read the meta-analysis at its stated strength: minimal effect on motor function, which is not the same as none. Quetiapine remains the reasonable first attempt on tolerability grounds while not belonging to that recommended pair, and holding both statements at once is what a good answer demonstrates. Whether pimavanserin can be obtained in India is not established by any source used here, so clozapine is the half of the pair to teach as realisable.

PSYCHOTIC SYNDROME	RELATION TO THE PRIMARY ILLNESS	FIRST MOVE	ANTIPSYCHOTIC TREATMENT
Ictal and post-ictal psychosis	Tied to the seizures; self-remitting with supportive treatment	Prevent the seizures	No prophylactic antipsychotic; the Maudsley allows a short-term symptomatic option
Inter-ictal and para-ictal psychosis	Course similar to schizophrenia	Select for lowest seizure propensity	Often needed
Psychosis in Parkinson disease	Often visual hallucinations; anticholinergics and dopamine agonists implicated	Exclude delirium, optimise environment, reduce or stop the implicated drugs	Only after those steps, and then quetiapine before clozapine

15.8 Depression and cognition in Parkinson disease

Depression in Parkinson disease has a ladder of its own, and its top step is counter-intuitive. A 2022 network meta-analysis reported in the Maudsley guidelines showed that dopamine agonists

and SSRIs have the highest efficacy and acceptability, so the class that treats the movement disorder and the class that treats depression share the top of the ranking rather than competing for it.

Augmentation with dopamine agonists such as pramipexole may be considered, but these drugs increase the risk of **impulse control disorders** and have rarely been associated with the development of psychosis. That is the same drug group implicated in inducing psychosis a moment ago, so the antidepressant strategy and the antipsychotic strategy pull against each other in one patient, and which way the prescription moves depends on which syndrome is currently dominant. The permissive modal is worth preserving in an answer: this is an option that may be considered, not a recommended step.

Electroconvulsive therapy sits further down the same ladder, and its risk signature here is not the one it had in epilepsy. Depression and motor symptoms generally respond well, but the risk of inducing delirium is high, particularly in patients with pre-existing cognitive impairment. One modality, two neurological illnesses, two opposite cautions: in unstable epilepsy its anticonvulsive property is the argument for it, and in Parkinson disease with cognitive impairment the delirium risk is the argument for care.

Cholinesterase inhibitors close the picture and are stated at a strength candidates routinely overstate. A Cochrane review and some large randomised trials concluded that there is evidence that they lead to improvements in global functioning, cognition, behavioural disturbance and activities of daily living in Parkinson disease, with the improvements in cognitive functioning described as modest, while motor function may deteriorate with a particular increase in tremor. Read it at its stated strength: a claim that evidence exists, hedged twice, with a trade of cognition against tremor that has to be discussed with the patient rather than resolved on their behalf.

> In epilepsy, treat the psychiatric illness anyway and choose the agent that does not provoke the seizure; in Parkinson disease, subtract the dopaminergic drug before adding an antipsychotic, and if one is still needed, quetiapine before clozapine.

16 · Psychopharmacology in pregnancy, lactation, the elderly and children

The drug does not change; the body around it does. Pregnancy, lactation, old age and childhood each alter clearance, distribution and vulnerability, and each adds a second person, or a failing organ, to a decision otherwise about one patient. The marks are not for reciting which agent is safe but for the decision under the changed pharmacology: which harm is traded against which, when the trade has to be made, what is measured afterwards, and which figures the standard sources dispute.

16.1 The decision is made before conception, not at the positive test

Organogenesis does not wait for a pregnancy test. The Maudsley Prescribing Guidelines place the period of maximum risk to the fetus from lithium at 2 to 6 weeks after conception, before many women know they are pregnant, and add that the risk of atrial and ventricular septal defects may

also be increased. The Indian Psychiatric Society bipolar guideline draws the operational conclusion: most pregnancies are unplanned or unrecognised in the early weeks, so contraceptive counselling is offered to all reproductive-age women with bipolar disorder on treatment. A conversation that begins when the test turns positive has missed the window it was meant to influence.

Folate sharpens the same point. The Maudsley states there is no evidence folate protects against anticonvulsant-induced **neural tube defects** if given during pregnancy, but that it may do so if given prior to conception, the neural tube being essentially formed by 8 weeks of pregnancy; folate may nonetheless benefit early neurodevelopment and so should always be offered. The Indian bipolar guideline hedges the protective claim harder while keeping the same timing. The timing rule is firm; the protective effect is not.

RULE **Relapse is an exposure too.** The Indian schizophrenia guideline states that an effective antipsychotic on which the patient has been stable should not be discontinued or changed, because relapse poses a greater risk to the mother and baby than the potential adverse effects. Every decision here compares two harms rather than choosing between a harm and safety; an answer that lists teratogenic risks without the untreated illness on the other side of the scale has answered half the question.

16.2 Comparative teratogenic risk, and what the comparison licenses

Substitution rests on the spread between agents rather than on any single figure, so the risk of **major congenital malformations** is best held as a comparison across the drugs actually chosen between.

AGENT	MALFORMATION RISK AS THE SOURCE STATES IT	WHAT THE SOURCE DOES WITH IT
Lithium	Ebstein's anomaly in **0.05 to 0.1 per cent of live births** after first-trimester exposure, about twenty times the general population rate, as the Indian guideline states it. That Ebstein-specific multiple is the historical estimate; the modern cohort figure sits below the table	Small in absolute terms, which is why lithium may still be continued after shared decision-making
Valproate and divalproex	8 to 10 per cent, including neural tube defects	Contraindicated in pregnancy, and linked to later neurodevelopmental disorders
Carbamazepine	3 to 6 per cent	Adds a **neonatal bleeding diathesis** from vitamin K deficiency
Oxcarbazepine	Comparable to the general population rate, at approximately 3 per cent	Can be used in pregnancy, although safety data are limited
Lamotrigine	Approximately 3 per cent, similar to baseline, with no definite evidence of neurobehavioural sequelae	Preferred where the course is depression-predominant; clearance rises, so the dose often has to follow
Second-generation antipsychotics	4.5 per cent in exposed against 3.5 per cent in unexposed individuals	No significant increase after adjustment, apart from a small risperidone signal a later large study did not replicate
Selective serotonin reuptake inhibitors	Small increased absolute risk of cardiac defects with paroxetine, 2 per 1000 exposed pregnancies, which is an excess and not the total incidence	Generally safe if indicated; avoid paroxetine

Two readings are examinable. First, the spread is the argument: because one agent sits several times above its nearest comparator, the regulatory response is specific to that agent rather than a class-wide caution, and a within-class alternative at baseline risk is what makes a restriction workable instead of merely punitive. Second, absolute and relative risk are different instruments pointed at one thing. The Maudsley guidance on advanced prescribing states the lithium risk as an absolute frequency: a population risk of Ebstein's anomaly of 1 in 20,000, rising with lithium exposure in pregnancy to 1 in 1000 to 2000. That is the Indian figure in another currency rather than a competing estimate, and both are the same historical Ebstein-specific estimate. Contemporary cohort evidence has moved the magnitude. Patorno and colleagues, reporting a large registry cohort in the New England Journal of Medicine in 2017, found cardiac malformations overall in 2.41 per cent of lithium-exposed pregnancies against 1.15 per cent of unexposed, an adjusted risk ratio of 1.65 that rose with dose, and did not establish a twenty-fold

Ebstein-specific excess. Quote the guideline figure as the guideline's, and know that the twenty-fold multiple is the older estimate rather than the current one.

16.3 Continuing lithium: the scan, and two monitoring cadences

Continuing lithium commits the clinician to fetal surveillance and to a level schedule, and both are specified differently by different bodies. The Indian bipolar guideline states that in all pregnancies with first-trimester lithium exposure a 20-week fetal cardiac anomaly scan is recommended. The reproductive-safety table of that same document advises fetal echocardiography over a window rather than at a single time point, and this chapter does not hold that second figure. The disagreement is internal to one guideline and should be quoted as such.

Haemodynamic change in pregnancy alters lithium handling, and the two guidelines specifying a cadence do not specify the same one. The Indian bipolar guideline gives monthly monitoring till 36 weeks, weekly from 36 to 40 weeks, and twice each week in the first 2 weeks postpartum, framed as a suggestion and not a mandate. The Maudsley reports that NICE recommends lithium levels be checked every four weeks in pregnancy and weekly after 36 weeks. Print both, attributed. The Indian schedule carries a postpartum limb the United Kingdom summary does not reproduce, and that absence is a silence in the summary rather than a statement that no postpartum monitoring is specified.

16.4 Valproate: a prohibition, and the machinery that enforces it

Valproate is the one agent here where the answer is not a trade-off: the Indian bipolar guideline records it as contraindicated in pregnancy, and the figure above is the whole of the reason. The regulatory machinery is examinable in its own right. The Maudsley compresses the United Kingdom position into one line: valproate may not be initiated in patients under 55 or continued in women of child-bearing potential unless two specialists agree and document that there is no other effective or tolerated treatment. That compression welds two different rules together and drops an exception, and the regulator separates them. Under the Medicines and Healthcare products Regulatory Agency measures of January 2024, valproate must not be started in a new patient of either sex under 55 unless two specialists independently consider and document that there is no other effective or tolerated treatment, or that there are compelling reasons the reproductive risks do not apply. Continuation in a woman of child-bearing potential is governed instead by the pregnancy prevention programme, with annual specialist review and a signed risk acknowledgement, and not by the two-specialist initiation test. The Indian bipolar guideline quotes that fuller regulatory text, second escape clause included, which the one-line summary omits. Alongside the two-specialist rule sits the programme requirement: valproate should not be offered to girls or young women of child-bearing potential unless a **pregnancy prevention programme** is in place.

That programme is a United Kingdom instrument. No Indian counterpart is established by the sources used here: the Indian guideline refers only obliquely to restrictions on divalproex in reproductively active women and men, naming no programme, no contraception requirement, no

age bar and no two-specialist rule. Name the jurisdiction and say the equivalent is not documented.

The stopping decision has a defined shape. Where a woman on valproate decides to conceive, the Australian perinatal guidance reported by the Maudsley is to stop valproate over 2 to 4 weeks while adding high-dose folic acid **5 mg/day**, continued through the first trimester, and the Indian bipolar guideline corroborates that supplementation dose for the pre-conception period. That supplementation figure is a daily amount rather than a dose to build up to: no initiation step sits below it in either document, and neither states a maximum above it for preconception use.

GOOD TO REMEMBER

Where carbamazepine is carried to term, the vitamin K limb is what gets forgotten, because it is not a psychiatric act. What is established is the neonatal limb: intramuscular vitamin K for the newborn at birth, arranged with the paediatric team. Routine antenatal vitamin K for the mother purely because she is taking an enzyme-inducing anticonvulsant is not supported by adequate evidence, and the Royal College of Obstetricians and Gynaecologists green-top guidance on epilepsy in pregnancy makes maternal treatment a response to a demonstrated coagulation abnormality or another obstetric indication rather than a routine. The malformation risk was settled months earlier; the neonatal bleeding diathesis is still preventable at delivery.

16.5 Antipsychotics and antidepressants: the dose that should not fall

The commonest prescribing error in pregnancy is a well-intentioned dose reduction. The Indian schizophrenia guideline states that the dose should not be reduced merely because a patient is pregnant, since antipsychotic metabolism may increase during pregnancy and lead to subtherapeutic levels, and names clozapine as the exception, whose levels may instead increase by the third trimester. The direction of the kinetic change decides the direction of the dose change.

PITFALL

Halving an antipsychotic at the booking visit, reasoning that less drug means less exposure. Metabolism may already be rising for the agent in question, in which case the reduction compounds a fall that was happening anyway. The rule is not that doses go up in pregnancy. The American College of Obstetricians and Gynecologists puts it as do not reduce automatically, and change a dose only where that drug's pharmacokinetics, the clinical response or a level justify the change. Where a dose was raised during pregnancy, the postpartum step is to bring it back down under close monitoring rather than leave it where the pregnancy put it.

Antidepressant decisions turn on two late-pregnancy risks rather than on malformation. The Indian guideline for obsessive-compulsive disorder states that selective serotonin reuptake inhibitors taken late in pregnancy may increase the risk of **persistent pulmonary hypertension of the newborn** by more than twofold, giving the absolute risks as 3 infants per 1000 exposed against 1.2 per 1000 unexposed. The Indian bipolar guideline describes that same association only as conflicting evidence in the third trimester, so even the strength of the claim is unsettled within one national guideline set. Separately, **neonatal adaptation difficulties** may occur in

about one-third of infants exposed late in pregnancy, which makes a paediatric handover the response rather than stopping the drug before delivery.

TRAP **One paroxetine risk, two incompatible figures inside the same national guideline set.** The bipolar guideline gives the small increased absolute risk of cardiac defects as 2 per 1000 and instructs that paroxetine be avoided. The guideline for obsessive-compulsive disorder, from the same series, states the septal-defect risk as a percentage an order of magnitude higher against a stated population rate, adding that the evidence is inconsistent. This chapter holds only the per-thousand figure. Quote what you hold, name the source, and say the other document disagrees.

16.6 Lactation: the pregnancy answer does not carry across

The reversal between the two phases is the teachable point. The Indian bipolar guideline states, in a single sentence, that breastfeeding is generally avoided with lithium because of the risk of neonatal toxicity and the need for infant serum drug level monitoring, and that divalproex is considered relatively safer in breastfeeding. Read it carefully: the comparison between the two agents is carried by a semicolon rather than asserted, and the agent contraindicated in pregnancy is the one treated more permissively at the breast. The next sentence supplies the conditional exception, that breastfeeding during lithium treatment may be considered in healthy, full-term infants with close monitoring, if only because discontinuation may pose greater risks. Emergency management of established lithium toxicity belongs to the Perinatal/women's mental health and emergency psychiatry notes.

Clozapine is where two Indian guidelines part company outright. The bipolar guideline states that clozapine should not be prescribed during breastfeeding, because of rare but serious adverse effects warranting blood monitoring. The schizophrenia guideline's corresponding table states only that antipsychotics are not contraindicated during breastfeeding and that clozapine may be present in breast milk, its potential risks to the infant needing to be considered and monitored. Both belong in an answer. For the remaining second-generation agents the bipolar guideline states that current evidence suggests quetiapine and olanzapine are safe during breastfeeding, with lactation data somewhat limited for the others.

Among antidepressants, the guideline for obsessive-compulsive disorder separates the SSRIs by how much of the drug reaches a breast-fed infant, placing sertraline, fluvoxamine and paroxetine at the low end and fluoxetine and citalopram above them, and it then surmises that they are relatively safe during breastfeeding, the hedge being the guideline's own. The two percentages that carry that separation are deliberately not printed here, because the sentence holding them describes them at once as concentrations in the plasma of breast-fed infants and as a share of the maternal dose, and those are two different quantities with two different denominators, since a plasma concentration cannot be a percentage of a dose. The source is internally inconsistent at the exact point a candidate would quote it. Reading the figures as relative infant dose, the share of the maternal dose reaching the infant through milk, is the most plausible repair, but it is an interpretation and not what the guideline says, and no source read here gives a drug-specific relative infant dose or an infant plasma ratio to substitute for it. Nothing is printed on either

reading. Milk supply is a separate axis from infant exposure, and it is where an Indian dataset displaces a Western default: **complete lactation failure** is a documented aripiprazole effect.

TRAP **Attaching a percentage to the wrong denominator.** An Indian perinatal-psychiatry chart review found that 60 of the 398 women attending the service were prescribed aripiprazole during pregnancy, with lactation data available for the 35 who continued the drug during the postpartum period; of those 35, 26 (74%) experienced complete lactation failure and 4 (11%) had insufficient milk production. The denominator of both percentages is the thirty-five with lactation data, not the sixty prescribed the drug. A national guideline compresses the study into a retrospective analysis of sixty women, attaching both percentages to the wrong number, so a candidate quoting the guideline sentence inherits an error the primary paper does not contain.

IN INDIA

Aripiprazole is the Western default for metabolic safety, and Indian perinatal data attach a lactation penalty to it that Western reproductive-safety tables do not show. A woman established on aripiprazole who intends to breastfeed therefore needs her postpartum plan settled antenatally, while there is still time to change the agent, rather than after the milk has failed. Where an antidepressant is needed alongside, the Indian bipolar guideline names sertraline as having the most substantial evidence of lactation safety.

MNEMONIC

Plan, Prefer, Persist, Postpartum

The perinatal spine. Plan before conception, because the highest-risk weeks precede the test. Prefer the agent whose documented risk is lowest for this illness course. Persist with an effective drug rather than reduce it because the physiology changed. Re-plan at delivery, when clearance, feeding and the level schedule all reverse together.

16.7 The elderly: assume the reserve has already gone

Age is a kinetic diagnosis before it is a clinical one. The Maudsley quantifies the decline that governs every renally cleared psychotropic.

35%	50%	two-thirds
RENAL FUNCTION LOST BY SIXTY-FIVE	LOST BY EIGHTY	FUNCTION TO ASSUME AT MOST

The third figure is the Maudsley's default assumption rather than a measurement, and it is a place to begin thinking rather than a dosing method. Renal function in an older adult is estimated before a renally cleared psychotropic is dosed, not assigned from age: the KDIGO chronic kidney disease guideline of 2024 makes estimated glomerular filtration rate or creatinine clearance the instrument, and adds that a creatinine-based estimate can overstate function where muscle mass is low, which is common in exactly this population, so cystatin C or a combined equation is used where creatinine is unreliable. The underlying pharmacokinetic principles belong to the Psychopharmacology I: principles and core drug classes notes. For maintenance treatment in bipolar disorder the Indian guideline's suggested serum lithium levels are 0.6 to 0.8 mEq/L in

medically healthy individuals and **0.4 to 0.6 mEq/L in older adults or those with relevant medical issues**. The acute-phase figures in that guideline do not agree between its dosing table and its running text, and neither set is printed here; target concentrations and sampling timing are taught with therapeutic drug monitoring. The standard texts also diverge on the elderly starting dose and on how soon the first level is drawn, and no figure for either is printed.

For depression in older people the Indian guideline directs start low and go slow, with monitoring especially for hyponatraemia and falls, and states that older patients usually require half the starting and target dose used in adults. The Indian schizophrenia guideline sets a different, smaller reduction for antipsychotics, and the two must not be merged into one rule. Choice of agent is then driven by two burdens rather than by efficacy: that guideline states that first-line antipsychotics in older adults will be the ones with lower **anticholinergic burden** and lower sedative property, to prevent falls, that anticholinergic co-prescription should be avoided to the possible extent, and that monitoring for metabolic adverse effects should be more frequent. The interval is nowhere quantified, and supplying one would be worse than leaving it open.

Antipsychotics in dementia carry a licence and a warning that must be quoted together. Risperidone is specifically indicated for short-term treatment, up to 6 weeks, of persistent aggression in moderate to severe Alzheimer's disease unresponsive to non-pharmacological approaches, and is licensed up to 1 mg twice daily although the optimal dose in dementia is 500 micrograms twice daily. Against that, the United States Food and Drug Administration's 2005 advisory determined an increased risk of death in older adults with dementia-related psychosis treated with atypical antipsychotic drugs of between 1.6 and 1.7 times that seen in placebo-treated patients. That estimate is pooled across placebo-controlled trials of several atypical agents, and it is the basis of a class boxed warning which risperidone carries; a standard textbook restates the same ratio as risperidone-specific, and the advisory does not support that reading. What the licence entry does not give is a dose to start at. It states the licensed upper figure and the optimal figure for dementia and nothing beneath either, so the optimal dose is where the source lands rather than a rung arrived at by titration. Prescribe, review at the stated horizon, stop.

16.8 Children and adolescents: the monitoring is the answer

Paediatric prescribing is examined through monitoring far more than through drug choice, because the choice is narrow and the surveillance specific. For depression, the Maudsley makes fluoxetine the recommended first-line medication in children and adolescents, started at 10 mg daily and increased after one week to the minimum therapeutic dose of **20 mg daily**. The entry stops there. It gives the start and the step and names no ceiling for this population, so the therapeutic minimum is a floor and must not be read backwards as the limit. Valproate reappears with a paediatric starting dose and the same prohibition attached: the Maudsley table of starting doses in children and adolescents gives valproate a weight-based starting range in divided doses, with plasma levels used to determine the maintenance dose, repeating in the same cell that valproate is not to be offered to girls or young women of child-bearing potential unless a pregnancy prevention programme is in place. The range itself is withheld here: that table row

carries no indication, and a paediatric weight-based dose quoted without the condition it treats is the shape of a wrong answer. A footnote weight-bands it, the lower end belonging to lighter children; the weight cut-off is not stated here, and the Indian adult figure for divalproex belongs to a different population. No ceiling appears in the valproate row either. It carries a starting range and a prohibition and nothing above them, the maintenance dose being fixed by plasma level rather than by any stated milligram maximum.

Stimulant monitoring is where two Indian sources diverge. The Indian Psychiatric Society guideline on attention-deficit hyperactivity disorder directs that heart rate and blood pressure be measured at every visit during titration and once every 3 months. An Indian postgraduate text gives a looser schedule: pulse, blood pressure and heart rate at each dose increase and every 6 months, with weight and height on standardized charts at 6-month intervals as a minimum, adding that routine haematological tests are unnecessary. The guideline's intervals are the ones to work to, and the textbook's looser schedule is worth knowing as the divergence rather than quoting as the standard. That guideline also contradicts itself internally, its continued-monitoring text asking for measurement at every visit with no step-down. Whether an electrocardiogram is required before a stimulant cannot be settled from these sources either: the textbook and the guideline do not agree on whether the tracing is a baseline requirement or a reactive step once a blood-pressure problem has emerged.

Two thresholds carry named consequences and both are frequently misplaced by tier. In the guideline's adverse-effect table, initial management of tachycardia and rise in blood pressure, listed for methylphenidate and atomoxetine, is triggered by sustained tachycardia of more than 120 beats per minute, or systolic or diastolic blood pressure more than the 95th percentile measured on 2 occasions, and the action is to reduce the dose and refer to a cardiologist. Separately, and only as a second-tier step once initial management has failed, persisting **growth suppression**, or growth parameters below the critical thresholds even after stopping the drug, roughly below the third percentile and growth velocity falling 2 centile lines over 1 year, prompts referral to a paediatric endocrinologist. That criterion is conjunctive: both limbs, not either alone. Initial management of growth retardation, in the order the guideline prints it, is a ladder.

- Dosage reduction.
- Weekend and longer drug holidays, to allow catch-up growth.
- Changing the class of medication.
- Stopping the drug.

The growth signal justifying all of this is real but modest. In the MTA follow-up, the group that did not receive medication was about 1 cm taller and 1 kg heavier at the 24-month follow-up than the group treated with stimulant medication for the entire two-year period. A difference of that size argues for charting growth at fixed intervals, not for withholding an effective treatment.

In every one of these populations the pharmacology of the drug is unchanged; what has changed is the body it enters, so the dose, the level and the monitoring interval must all move with it.

17 · QTc prolongation and cardiometabolic monitoring across psychotropics

Two monitoring problems share one heading for a reason. Cardiac conduction and cardiometabolic risk are the harms psychotropics produce slowly, silently and measurably, and that combination makes a schedule rather than a symptom the instrument of safety. Both are examined as decisions rather than as facts: which number obliges which action, which agent carries which liability, who gets a tracing and how often, and where an Indian guideline parts company with a British one. The receptor pharmacology behind each class's adverse-effect signature belongs to the Psychopharmacology I: principles and core drug classes notes. What follows is the monitoring that pharmacology obliges, and the decisions its results force.

17.1 The thresholds, and how firmly each is held

The **QTc** is the interval corrected for heart rate, and its published cut-offs are not held with equal confidence. The Maudsley 15th edition gives the normal limits as 440 ms for men and 470 ms for women and describes arrhythmia risk beyond those limits as exponentially related to the extent of prolongation, while labelling that limb very limited evidence with well-known exceptions. It holds the next step more firmly: values over **500 ms** carry a clearly increased risk of arrhythmia. The same passage carries a fourth limb, attaching a better-than-even chance of inducing torsades to a still longer interval, and that limb is withheld here with its figure: no QTc value has a validated patient-level probability of torsades attached to it, because the risk that a given tracing carries also depends on the drug, the correction formula used, the heart rate, the electrolytes and the structural state of the heart. Carry the strength of each statement along with its figure, because an answer that prints all three at one confidence has misreported the source it is quoting.

440 ms	470 ms	500 ms
NORMAL LIMIT, MEN	NORMAL LIMIT, WOMEN	CLEARLY RAISED RISK

A second Maudsley text puts coordinates on the same curve, and this is the point at which a printed figure has to be refused rather than quoted. The Maudsley Guidelines on Advanced Prescribing in Psychosis state that the relationship between QTc and torsade risk is exponential and attach an absolute probability to each of two named intervals. Those two percentages are withheld here. Cardiology guidance on the prevention of torsades de pointes in hospital settings describes the risk as rising continuously with the interval, substantially so above the clearly-raised-risk threshold already named, and declines to convert any single QTc value into a probability for an individual patient. So the two accounts differ in kind rather than in direction: one calls the exponential relation very limited evidence, one quantifies it, and the cardiology position is that the quantification does not survive at the level of the patient in front of you. Teach the direction and the action threshold, attribute whichever account you use to the body that states it, and do not read a percentage off a tracing.

17.2 Reading the number as an instruction

RULE **A QTc value is read as an instruction, not as a diagnosis.** Each band of the Maudsley management grid names two things at once, what happens to the prescription and how urgently cardiology is involved, and the grid carries a further row that fires on the shape of the tracing whatever the measured interval turns out to be.

Table 1.38 of the Maudsley Prescribing Guidelines sets out that ladder for patients receiving antipsychotic drugs. Read the direction of travel across it rather than the cells in isolation: the intermediate band leaves the clinician discretion over dose reduction or switching, and the top band withdraws that discretion and attaches an immediate referral to it.

QTC ON THE TRACING	WHAT HAPPENS TO THE PRESCRIPTION	CARDIOLOGY REFERRAL
Below 440 ms in men or 470 ms in women	None, unless T-wave morphology is abnormal	Consider if in doubt
Above those limits but below 500 ms	Consider reducing the dose or switching to a drug of lower effect, and repeat the electrocardiogram	Consider
Above 500 ms	Repeat the electrocardiogram, stop the suspected causative drug and switch to a drug of lower effect	Immediately
Abnormal T-wave morphology, whatever the interval	Review treatment, and consider reducing the dose or switching to a drug of lower effect	Immediately

That last row is the one candidates leave out, and it is the row that makes the grid a clinical instrument rather than an arithmetic one. **Abnormal T-wave morphology**, independent of the QTc value, triggers treatment review, consideration of dose reduction or switching, and immediate cardiologist referral. A comfortable interval on a tracing with an abnormal T wave does not discharge the clinician.

17.3 Three variables, not one: the drug, the co-prescription and the patient

The Maudsley stratification of antipsychotics by effect on the QTc runs to four bands, and the usable information sits at the two ends. The high-effect band carries pimozide and sertindole, while the no-effect band holds brexpiprazole, cariprazine, lurasidone and lumateperone. The low-effect band takes aripiprazole, asenapine, clozapine, olanzapine, paliperidone, risperidone, sulpiride and the older thioxanthenes and phenothiazines listed with them, and amisulpride, chlorpromazine, haloperidol, quetiapine and ziprasidone sit in the moderate band. The clinically useful reading is comparative rather than absolute: the choice is rarely between an agent that prolongs and one that does not, but between bands.

The co-prescription is the second variable, and it is usually written by somebody else. Table 1.37 of the same chapter groups the non-psychotropics associated with QT prolongation into four

families.

- Antibiotics: erythromycin, clarithromycin, co-trimoxazole and pentamidine, with a note that some quinolones affect the QTc. One further antibiotic printed on that row is withheld here, because the maintained database the same table points to does not carry it in any torsades-risk category and no causal source for it could be supplied.
- Antimalarials: chloroquine, mefloquine and quinine.
- Antiarrhythmics: quinidine, disopyramide, procainamide, sotalol, amiodarone and bretylium.
- Others: amantadine, cyclosporin, diphenhydramine, hydroxyzine, methadone, nicardipine and tamoxifen.

The table directs the reader to Crediblemeds.org for current information and states that it is not a complete list. That caveat teaches better than the list does: absence from a printed table is not evidence that an agent is safe, and the reference to a maintained database is the guideline conceding that a book cannot keep this current. The converse holds too, and it is why one entry was withheld above: presence on a printed row is not itself a risk category, and where the maintained database and the book disagree it is the database that the book tells you to use.

The patient is the third variable. The Maudsley Table 1.36 of physiological risk factors sorts them into cardiac factors, being **long QT syndrome**, bradycardia, ischaemic heart disease, myocarditis, myocardial infarction and left ventricular hypertrophy; metabolic factors, being **hypokalaemia**, hypomagnesaemia and hypocalcaemia; and others, being extreme physical exertion, stress or shock, anorexia nervosa, extremes of age and female gender. Of the three groups only the metabolic one is correctable by a laboratory result and a bag of fluid, which is why it is checked first.

GOOD TO REMEMBER

These additional risk factors seem almost always to be present in cases of antipsychotic-induced **torsade de pointes**. Treat that sentence as an instruction about where to look: a prolonged interval in a patient on an antipsychotic is a prompt to hunt for the second and third variables rather than a verdict on the first. Send the electrolytes, read the rest of the drug chart, and find out what the heart was like before the prescription was written.

17.4 Antidepressants, and one disagreement worth knowing

The antidepressant account is narrower than the antipsychotic one and easier to overstate. Among antidepressants, the Maudsley Prescribing Guidelines hold that QT changes are not usually seen with most SSRIs at normal clinical doses, that the class-level association is largely driven by citalopram and escitalopram and is dose-related but modest, that sertraline prolongs QT by 5 to 10 ms at 400 mg a day, a figure the same evidence sets against those normal clinical doses rather than a dose it recommends. A fourth clause, exonerating three named newer antidepressants of any effect on cardiac conduction, is withheld: the passage quoted from that chapter does not carry it, and an absolute of that width is not a thing to print on an unverified line. Treat the position on those agents as unstated here rather than as reassurance.

The Indian Psychiatric Society depression guideline declines the second of those clauses. It states that escitalopram does not pose a clinically significant risk of QTc prolongation, which is a softer position than attributing the class effect to citalopram and escitalopram jointly. Both belong in an answer: the Indian guideline exculpates escitalopram and the United Kingdom source does not, neither position is derivable from the other, and a candidate who quotes only one is wrong against the examiner who holds the other.

Citalopram, one of the antidepressants named in that class effect, is the standard pharmacovigilance illustration in this territory: the FDA restricts it to a maximum of **20 mg daily** for defined subgroups, including patients taking a CYP2C19 inhibitor and those older than 60 years, because it causes dose-dependent QT interval prolongation. That list is not closed, and reading it as closed is the error to avoid: the same 20 mg ceiling is set by the regulator's 2012 safety communication for patients with hepatic impairment and for CYP2C19 poor metabolisers. Note what kind of ceiling that is. It is not a toxicity limit and not a licensing quirk; it is a regulator converting a conduction signal, in subgroups defined by enzyme inhibition, enzyme genotype, liver function and age, into a hard dose restriction. Note also that the citalopram restriction travels in one direction only. It names a maximum for each subgroup and no starting dose at all, so it constrains where a regimen may end up and says nothing at all about where it starts.

17.5 Who gets a tracing, and when

Two bodies answer this differently and both answers are examinable. The Maudsley position is to measure QTc in all patients prescribed antipsychotics on admission and, if there is a previous abnormality or a known additional risk factor, at the annual physical health check. Its monitoring table places the electrocardiogram at baseline, at target dose, on admission to hospital, if cardiac symptoms occur, and when medication is changed; the tracing is mandatory for haloperidol, pimozide and sertindole, and for ziprasidone in some situations, and ideally all patients should be offered an electrocardiogram at least yearly.

The Indian Psychiatric Society schizophrenia guideline triggers on different criteria. For the chronic phase or multiepisode relapsing-remitting illness it states that an electrocardiogram would be needed for patients over 40 years, on injectable antipsychotics, or on antipsychotics at high risk of increasing QTc, along with nutritional status. Both documents use a drug trigger, so the contrast is narrower than it first looks. What is genuinely Indian is the age threshold, which the Maudsley triggers do not use, and the naming of the injectable route as a trigger in its own right.

Opioid substitution carries a third grid, and it is worth knowing because the same figures do different work in it. The Maudsley Prescribing Guidelines, which carry this table across successive editions with identical thresholds, grade methadone electrocardiogram monitoring by sex: borderline prolonged QTc is 470 ms or more in females and 440 ms or more in males, prolonged QTc is 500 ms or more, and very prolonged QTc is 550 ms or more, each band carrying its own action limb. The two numbers that were the upper edge of normal in the antipsychotic grid are the entry to the borderline band here, which is a reminder that a threshold belongs to the table it was printed in.

17.6 The cardiometabolic panel and the intervals that carry it

Metabolic monitoring is where the Western and Indian answers diverge furthest. The parameters are much the same in every grid; what differs is how early the first repeat falls, how the first year is treated, and which agents earn a column of their own. The accompanying figure sets the recurring intervals side by side.

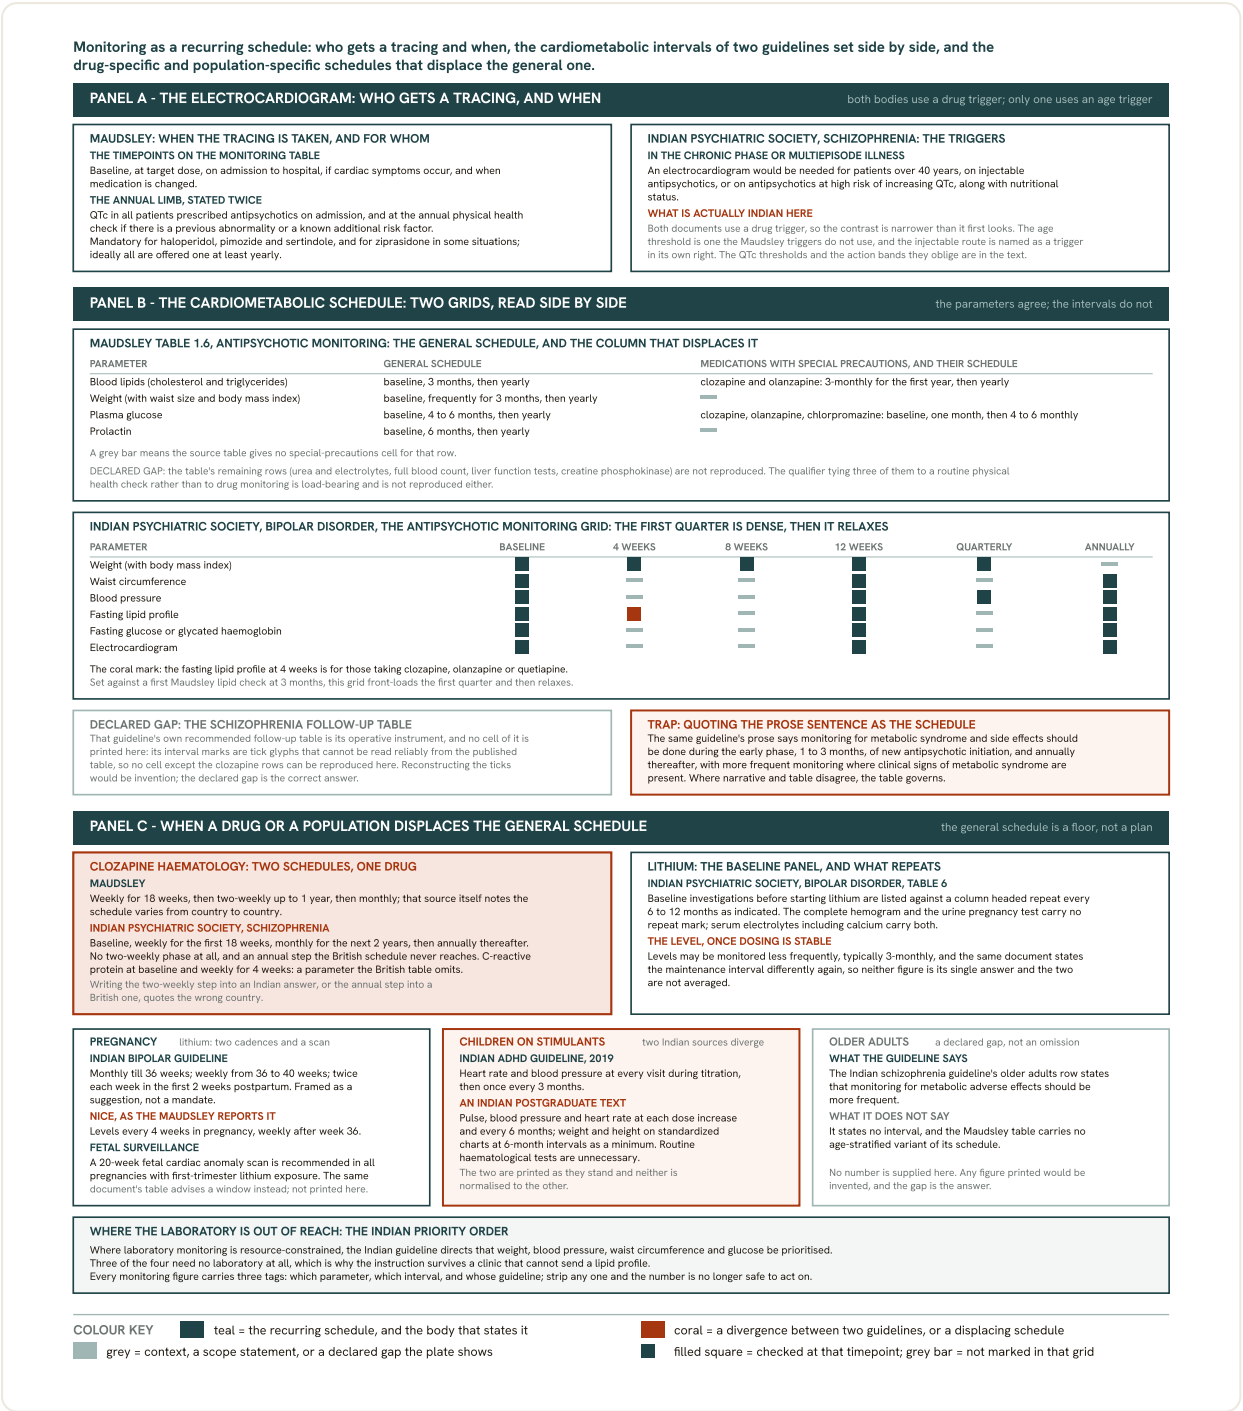


Fig 37.8 The recurring monitoring schedule: who gets a tracing and when, the cardiometabolic intervals of two guidelines set side by side, and the drug-specific and population-specific schedules that displace the general one. The Maudsley table is built around one early check and an annual rhythm, and carries a second column for clozapine, olanzapine and chlorpromazine that displaces the general one; the Indian bipolar grid front-loads the first quarter, with visits at 4, 8 and 12 weeks, and then relaxes. Clozapine haematology and lithium each carry a whole schedule of their own in place of the general one, and the two clozapine schedules part company at the second step: the British goes weekly, two-weekly, monthly, while the Indian goes weekly, then straight to monthly, and eventually to annual. Pregnancy, childhood and old age each set their own interval, and the interval for older adults is a declared gap, because the guideline says monitoring should be more frequent and states no number. The Indian schizophrenia guideline's own follow-up table is named as the operative instrument and is not reproduced, because its interval marks cannot be read reliably from the published table. Plasma target ranges are not on this plate; they are taught with therapeutic drug monitoring. (Maudsley Prescribing Guidelines in Psychiatry, 15th ed.; Indian Psychiatric Society Clinical Practice Guidelines 2025 for schizophrenia and bipolar disorder; for the stimulant row, the Indian Psychiatric Society attention-deficit hyperactivity disorder guideline of 2019, from its child and adolescent set, not the 2025 set; an Indian postgraduate child and adolescent psychiatry text; NICE as reported by the Maudsley.)

The Maudsley Table 1.6 gives blood lipids at baseline, at 3 months, then yearly; weight, including waist size and body mass index where possible, at baseline, frequently for the first 3 months, then yearly; plasma glucose at baseline, at 4 to 6 months, then yearly; and prolactin at baseline, at 6 months, then yearly. That is a schedule built around a single early check and an annual rhythm thereafter.

PITFALL

Applying the general interval to a patient who is on an agent with a column of its own. In the same Maudsley table the medications with special precautions on the blood lipids row are clozapine and olanzapine, tested 3-monthly for the first year, then yearly; on the plasma glucose row they are clozapine, olanzapine and chlorpromazine, tested at baseline, at one month, then 4 to 6 monthly. The general and the special-precautions schedules sit in adjacent columns of the same row, and reading across the wrong one is the commonest way a monitoring plan is written short for exactly the patients who needed it longest.

The Indian Psychiatric Society bipolar guideline's antipsychotic monitoring grid is denser early. Its columns are baseline, 4 weeks, 8 weeks, 12 weeks, quarterly and annually. Weight with body mass index is marked at baseline and at each of the three early visits and quarterly, but not annually; waist circumference at baseline, at 12 weeks and annually; blood pressure at baseline, at 12 weeks, quarterly and annually; the fasting lipid profile at baseline, at 12 weeks and annually, with an added check at 4 weeks for those taking clozapine, olanzapine or quetiapine; and fasting glucose or glycated haemoglobin and the electrocardiogram at baseline, at 12 weeks and annually. Set against a first Maudsley lipid check at three months, the Indian grid front-loads the first quarter and then relaxes.

TRAP **Quoting a guideline's prose sentence as though it were the guideline's schedule.** The Indian Psychiatric Society schizophrenia guideline states in its text that investigative monitoring for metabolic syndrome and side effects should be done during the early phase, 1 to 3 months of new antipsychotic initiation, and annually thereafter, with more frequent monitoring where clinical signs of metabolic syndrome are present. That sentence is not the document's operative instrument: the same guideline carries a recommended follow-up table, and where the narrative and the table do not agree the table governs. Name the prose statement, attribute it, and do not present its interval as the Indian metabolic monitoring schedule.

17.7 Clozapine haematology: two schedules, one drug

TRAP **Assuming one clozapine blood-count schedule worldwide.** The Maudsley schedule is weekly for 18 weeks, then two-weekly up to 1 year, then monthly, and that source itself notes the schedule varies from country to country. The Indian Psychiatric Society schizophrenia guideline runs baseline, weekly for the first 18 weeks, monthly for the next 2 years, then annually thereafter, with no two-weekly phase at all and an eventual annual step the British schedule never reaches. Writing the two-weekly step into an Indian answer, or the annual step into a British one, quotes the wrong country.

The action thresholds travel with whichever schedule is being used. The suspect medication is stopped if neutrophils fall below **1.5 x 10^9/L** in a patient taking clozapine, unless **benign ethnic**

neutropenia is diagnosed, and specialist medical referral is made if neutrophils fall below 0.5 x 10^9/L on clozapine. The exception is written into the rule rather than bolted onto it, so a low count is a question about the patient's baseline before it is an instruction to stop.

The Indian guideline adds a parameter the British table does not: C-reactive protein for clozapine at baseline and weekly for 4 weeks. Its own footnote is candid about the rest of the schedule, recording that an emerging expert statement suggests less stringent monthly blood monitoring after the first year of clozapine therapy because the incidence of clozapine-induced agranulocytosis is very low, but that this has yet to be reflected in any guidelines worldwide for want of robust evidence. The monthly figure therefore stands, and an answer that softens it on the strength of the footnote has read the footnote backwards.

17.8 Lithium's baseline panel, and the Indian adaptation

Lithium is monitored on two tracks and only one of them is a plasma concentration. The Indian Psychiatric Society bipolar guideline's Table 6 lists the baseline investigations to be done before starting lithium against a column headed repeat every 6 to 12 months as indicated: the complete hemogram and the urine pregnancy test carry no repeat mark, while serum electrolytes including calcium carry both. The same document also states that once lithium dosing has been stabilised the blood levels may be monitored less frequently, typically 3-monthly, and its narrative text states the lithium maintenance interval differently again, so neither figure is the guideline's single answer and the two must not be averaged into one.

What the Indian documents add beyond intervals is a definition and a priority order. The criteria for **metabolic syndrome** are population-specific, and the guideline gives its reason in the open: population-specific criteria should be followed when assessing metabolic syndrome, given the higher body fat content and risk profile in Indians.

IN INDIA

Read the list below as components, not as a checklist that must be completed. Under the harmonised definition of the metabolic syndrome, the diagnosis is made when any three of the five components are present, so a patient who meets three of them has the syndrome and a patient who meets two does not. Applied to an Indian patient, those five components are a **waist circumference** at or above **90 cm** in males and 80 cm in females, triglycerides at or above 150 mg/dL, HDL below 40 mg/dL in males and below 50 mg/dL in females, systolic blood pressure at or above 130 mmHg and/or diastolic at or above 85 mmHg, and fasting blood sugar at or above 100 mg/dL. Assess the patient in front of you against these, not against a threshold set imported from another population, because the same measurement can clear one definition and meet the other.

MNEMONIC

Weight, Pressure, Waist, Glucose

Where laboratory monitoring is resource-constrained, the Indian Psychiatric Society schizophrenia guideline directs that weight, blood pressure, waist circumference and glucose be prioritised. Three of the four need no laboratory at all, which is why the instruction survives a clinic that cannot send a lipid profile, and it answers a question the Western tables never ask: what to keep when the full panel is out of reach.

Every monitoring figure carries three tags: which parameter, which interval, and whose guideline; strip any one of them and the number is no longer safe to act on.

Neurostimulation and device-based treatments

18 · Electroconvulsive therapy: indications, mechanism, technique and adverse effects

Every other treatment in this chapter is a molecule. **Electroconvulsive therapy** is the one intervention chosen against the drugs rather than among them, and the long modality question asks the same five things each time: what it is, how it is thought to work, which disorders it earns a place in, how good the evidence is in each of them, and what it costs the patient. The procedural act itself, the technique and its parameters, is set out in the ECT and neurostimulation notes. What follows is the modality and its utility across disorders, which is where the marks are actually lost, usually by quoting one guideline body as though it spoke for all of them.

18.1 What is established about mechanism, and what is not

An Indian review of how the treatment works reaches a conclusion that is itself the examinable point: **no single mechanism of action** is established, because the cause and effect relationship between the observed biological findings and the therapeutic effect has not been settled, and a single explanatory account cannot prudently be assumed. Candidates lose marks by picking one of the competing hypotheses and defending it as the answer. The stronger answer sets them out as convergent lines of evidence and says plainly that none has been shown to carry the therapeutic effect on its own.

RULE **Seizure induction appears necessary but not sufficient for the antidepressant effect.** A generalised seizure can be induced and still fail to produce an antidepressant response, which constrains every purely seizure-based account of how the treatment works. It also explains why the modality is not interchangeable with any other way of producing a seizure, and why an answer that stops at "it induces a seizure" has answered nothing.

The integrative account now in circulation is explicitly **a proposed model** and not an established mechanism: an acute disruption of brain function, followed by a transient enhancement of neuroplasticity, with the combination of disruption and neuroplastic change rewiring the circuits implicated in depression. Write it as a model. Presenting it as settled mechanism is a correctness error, and it is the easiest one to make in this section because the model is coherent, widely

repeated and pleasant to remember. The useful discipline is to attach the evidential status to the sentence rather than to a footnote: unsettled mechanism, one plausible integrative model, one constraint that any model has to satisfy.

18.2 The indication set as the Indian guideline draws it

The Indian Psychiatric Society guideline on the use of electroconvulsive therapy sets the treatment as first-line where a rapid or definitive response is needed to avert harm to self or others, and names its exemplar situations: acute suicidal risk, agitation, **catatonia**, and deteriorating physical status secondary to a psychiatric condition. Two features of that sentence are examined. The trigger is the need for speed or certainty of response, not a diagnosis; and the guideline's own wording marks the list as open rather than closed, so it is an illustration of the trigger and not a set of four qualifying conditions.

MNEMONIC

Suicide risk, Agitation, Catatonia, Deterioration

The four situations the Indian guideline names when it places the treatment first-line. Recite the list, then immediately say that the guideline marks it non-exhaustive, because a candidate who presents four exemplars as four criteria has misread the recommendation in the direction that matters clinically.

Beyond that trigger, the same guideline enumerates indications disorder by disorder in a table, and the depression row is the spine of the Indian indication list. That row is reproduced in the comparison below rather than here. What the reader should take from the guideline as a document is its method: it was produced by consensus, with no formal systematic literature search. That is not a criticism, and it does not weaken the Indian practice position, which is what the document exists to state. It does decide where each half of an answer comes from. Effect sizes and response rates come from trials and meta-analyses; the position on when the treatment is offered in Indian practice comes from the guideline. Mixing the two, and citing the guideline for an effect size, is a sourcing error a good examiner will notice.

18.3 What each body recommends, side by side

The bodies do not say the same thing, and the divergence is not error: they are answering different questions about different populations, with different methods behind them. Reading the recommendations as one recommendation is the commonest way this material goes wrong.

ISSUING BODY AND RECOMMENDATION	POPULATION IT SPEAKS TO	POSITION, AND THE STRENGTH ATTACHED TO IT
Indian Psychiatric Society guideline on the use of electroconvulsive therapy	Any psychiatric presentation needing a rapid or definitive response to avert harm	First-line, with acute suicidal risk, agitation, catatonia and deteriorating physical status given as exemplars and the list left explicitly open
The same guideline, indication table	Major depressive disorder	Poor oral intake; high suicidal risk; high distress requiring rapid symptom remission; psychotic features; melancholic features; peripartum depression; treatment resistance; severe mixed affective features
NICE TA59 recommendation 1.1	Catatonia, and a prolonged or severe manic episode	Recommended only to achieve rapid and short-term improvement of severe symptoms, and only after an adequate trial of other treatment options has proven ineffective and/or where the condition is considered potentially life-threatening
NICE, on schizophrenia	Schizophrenia, general use	The state of the evidence does not allow general use to be recommended, which is a statement about evidence and not a safety prohibition
NICE NG222 recommendation 1.13.9	People whose depression has responded to a course	Start or continue antidepressant medication or a psychological intervention to prevent relapse; consider lithium augmentation of antidepressant medication
VA/DoD major depressive disorder guideline	Severe major depressive disorder with catatonia, psychotic depression, or need for a rapid, definitive response on medical or psychiatric grounds	Recommendation 20, a strong recommendation to offer the treatment with or without psychotherapy, on evidence the guideline itself rates very low

The last row repays a second look, because it is the cleanest available illustration of a distinction candidates blur. A strong recommendation resting on evidence the issuing body grades very low is not an inconsistency: strength encodes how confident the panel is that the benefits outweigh the harms for that population, and it can be driven by the severity of the untreated condition rather than by trial quality. The three limbs of the older NICE appraisal work the other way, and they are conjunctive rather than alternative: the aim has to be rapid and short-term improvement of severe symptoms, other options have to have failed or the condition has to be potentially life-threatening, and the patient has to have catatonia or a prolonged or severe manic episode. Rapid improvement on its own is not the threshold there.

TRAP **Citing the old technology appraisal as the current position in depression.** Its depression recommendations were superseded by the later depression guideline; only the catatonia, mania and schizophrenia limbs survive. A candidate who quotes the appraisal for depression is quoting a withdrawn recommendation, and the marks go to whoever names the right document for the right disorder rather than to whoever remembers the most recommendation numbers.

18.4 How strong the evidence is, disorder by disorder

Efficacy is not uniform across the indication list, and the evidence supporting each indication is of very different quality. The pooled analysis that remains the standard efficacy citation compared real against **simulated treatment** in depressive disorders, and separately against pharmacotherapy, in a total of eighteen head-to-head drug comparisons.

-0.91	-0.80	50%	0.55
VS SIMULATED, SES	VS DRUGS, SES	CLOZAPINE AUGMENTATION	MEMORY LOSS, SMD

The comparison against simulated treatment rests on six trials and 256 patients, with a confidence interval running from -1.27 to -0.54; the comparison against drug treatment rests on eighteen trials and 1144 participants, interval -1.29 to -0.29. Both favour the treatment significantly. Neither is a sequencing instruction. A head-to-head superiority result over pharmacotherapy is an argument against reflexively exhausting drug trials first, not an argument for reserving the modality until they are exhausted, and reading it the second way inverts what the numbers show.

In catatonia the reported figures are the highest in psychiatry and the weakest in design. Retrospective observational studies report response rates ranging from **80 percent to 100 percent**, from case series rather than randomised trials. The same review supplies the counterweight, and it has to travel with the figure: although pharmacotherapy had failed in 85 percent of the patients, the treatment improved 59 percent of the cases. Quoting the higher range alone overstates the indication.

In schizophrenia the evidence separates into two questions that are routinely conflated. For **clozapine-resistant schizophrenia**, a single-blind randomised trial found half of those given augmentation met the response criterion while none of those continuing clozapine alone did, with 47 percent responding in the crossover phase and an intention-to-treat sample numbering 39; the same report states that further research is needed on whether the improvement persists and whether maintenance treatment is required. For **treatment-resistant schizophrenia** in general, a systematic review found moderate-quality evidence of a positive effect on medium-term clinical response relative to standard care, with no clear and convincing advantage or disadvantage for other outcomes, and 14 of its 15 included studies at high risk of bias. Treating the augmentation trial as evidence for the general resistant population, or the review as evidence for the clozapine-resistant subgroup, is a correctness error in either direction.

18.5 Where it enters against the alternatives, and what has to follow it

The modality question is really a sequencing question. Each of the device-based treatments enters the pathway at a different point and for a different reason, and the accompanying figure sets all of them against the disorders they are used in and the strength of evidence behind each. What distinguishes this one is that it is the only modality in that frame with a randomised superiority result against drug treatment in depression, and the only one whose indication list is written around urgency rather than around failure.

Neurostimulation utility across disorders: all six device-based modalities in one frame, placed by what must have failed before each is offered, by the indications it has earned, and by the strength of the statement that grants each one.

SIX MODALITIES, THREE QUESTIONS: WHERE IT ENTERS, WHAT IT HAS EARNED, WHERE IT FAILS — read the last column as hard as the one before it

MODALITY and its physical logic	WHERE IT ENTERS, AND WHOSE STATEMENT SETS IT prior failure is counted in a different unit on every row	THE INDICATIONS IT HAS EARNED, AND ITS GRADE the grades are not interchangeable across disorders	NEGATIVE, NULL OR UNGRADED the rungs candidates leave out
ECT convulsive	IPS ECT guideline 2023: FIRST-LINE where a rapid or definitive response is needed to avert harm. Exemplars: acute suicidal risk, agitation, catatonia, deteriorating physical status; the list is explicitly non-exhaustive. Consensus, with no formal systematic literature search. NICE TA59 rec 1.1 binds THREE LIMBS TOGETHER: rapid and short-term improvement of severe symptoms; other options failed and/or the condition potentially life-threatening; and catatonia, or a prolonged or severe manic episode.	DEPRESSION Real against simulated ECT: SES -0.91, from 6 trials and 256 patients. Against pharmacotherapy: SES -0.80, from 18 trials and 1144 participants. VA/DoD rec 20 is strong FOR, on evidence the panel itself rated very low, in severe MDD with catatonia, psychotic depression, or a need for a rapid, definitive response. SCHIZOPHRENIA Clozapine-resistant: response 50% with ECT augmentation against 0% alone, 47% on crossover, ITT n=39. CATATONIA 80% to 100%, retrospective observational only.	Treatment-resistant schizophrenia generally: only moderate-quality evidence of medium-term response, with 14 of 15 studies at high risk of bias. The same catatonia review: pharmacotherapy had failed in 85% and ECT improved 59% of cases. NICE does not recommend general use in schizophrenia, on the state of the evidence, not on safety. TA59's depression limb was replaced by NG222.
rTMS sub-convulsive	IPS rTMS guideline 2023: an ADJUNCT to conventional treatment across the conditions reviewed, and in TRD augmentation beats stand-alone rTMS. VA/DoD rec 17: after partial or no response to two or more adequate pharmacologic trials. Weak, very low. RANZCP 2020 item 1.6: a reasonable number of adequate trials of BOTH drug AND psychological treatment, with no number fixed. Graded CBR, consensus and not evidence.	DEPRESSION graded STRONG: effect sizes 0.302 to 0.83, response odds ratios 3.26 to 3.64, remission 2.45 to 4.63. BIPOLAR DEPRESSION moderate: effect size 0.302, OR 2.72. SCHIZOPHRENIA negative symptoms SMD 0.49 to 0.64: high-frequency left DLPFC and more than 10 sessions beat sham. Resistant auditory hallucinations SMD 0.24 to 0.51, over 4 to 40 sessions given up to 8 weeks. OCD moderately positive: Y-BOCS SMDs 0.501 to 0.79, over 10 to 30 sessions across 1 to 12 weeks. GAD effect sizes are 1.45 to 1.87, with comorbid depression at SMD 1.65.	Panic disorder: not found effective. Clozapine-refractory schizophrenia: three RCTs of augmentation, no significant effect on positive or negative symptoms or on cognition. Tobacco use disorder is positive for craving and consumption; other substance use is not consistent. NOT recommended in acute suicidality. NICE: may be used with NORMAL arrangements for clinical governance and audit. MOST effect sizes above are for rTMS added to something else; verify each estimate.
tDCS subthreshold	No prior-failure count is set by any source read here. NICE puts tDCS for depression under SPECIAL arrangements for clinical governance, consent, and audit or research, one step below the position it gives rTMS. No Indian guideline sets it a place in the ladder at all.	DEPRESSION is the only disorder with a stated position, and both named trials are null on their own question. ELECT-TDCS: non-inferiority to escitalopram NOT shown over 10 weeks, with more adverse events; a single-centre trial. DepressionDC: the main trial showed no beneficial anti-depressant effect of active tDCS over sham.	No stated position in any other disorder. The mechanism is subthreshold polarisation, so the anode-excites heuristic cannot be written as a law: mechanisms vary with the current and its application.
IMPLANTED DEVICES DBS and VNS	DBS in OCD, US FDA HDE H050003, 10 February 2009: adults with chronic severe treatment-resistant OCD who failed AT LEAST THREE SSRIs, adjunct to medication. Humanitarian pathway on a probable-benefit standard, NOT effectiveness. VNS in depression, US FDA PMA P970003/S050, 15 July 2005: FOUR OR MORE adequate antidepressant treatments, at 18 or older, in a major depressive episode, chronic or recurrent. Eight guidelines: DBS only after all options have failed.	OCD carries the DBS regulatory position and DEPRESSION the VNS one, and neither rests on a positive controlled trial. DBS in depression: BROADEN found no significant difference between active and sham at six months and was halted on an interim futility analysis. VNS: acute D-02 responders 15.3% (17 of 111) active against 10.0% (11 of 110) sham, p=0.238; the 12-month uncontrolled phase gave response 29.8% and remission 17.1%.	VA/DoD recommends AGAINST DBS for depression outside research at strong strength, and suggests against VNS outside research at weak strength. NICE puts implanted VNS under special arrangements. One guideline in that same review does not recommend DBS beyond a research setting. RECOVER, the largest neuromodulation trial in psychiatry to date, did not meet its primary endpoint.
MST convulsive	Placed on no prior-failure ladder by any source read here. The comparator in both randomised trials is ECT, because the claim under test is not that seizures help but that this way of producing them costs the patient less.	CREST-MST, a confirmatory randomised non-inferiority trial, 239 randomised: remission 22.5% with MST against 27.8% with ultrabrief ECT, a 5.3% difference favouring ECT, with non-inferiority established; worsening of autobiographical memory 2.7% against 17.3%. BIPOLAR DEPRESSION CORRECT-BD pilot, n=55 randomised: remission 5/25 (20%) against 6/20 (30%), memory worsening 2/28 (7.1%) against 6/27 (22.2%). Not powered for efficacy.	Non-inferiority is not equality: the remission point estimate favoured ECT in both studies. The claim that MST seizures are less robust than ECT seizures rests on animal evidence only. No source read here states the disorder or population of the confirmatory trial, so this frame names none; only the pilot's own report names bipolar depression.
THETA-BURST AND ACCELERATED patterned rTMS	Placed on no prior-failure ladder by any source read here. VA/DoD rec 18 declines to recommend either for or against theta-burst stimulation in major depressive disorder, on insufficient evidence of very low certainty.	DEPRESSION THREE-D: iTBS non-inferior to 10 Hz rTMS in treatment-resistant depression, adjusted difference 0.103, lower 95% CI -1.16, against a pre-specified margin of 2.25. Stanford Neuromodulation Therapy: MADRS reduction 52.5% active against 11.1% sham at 4 weeks, 29 treated of 32 enrolled, at a planned interim analysis. No stated position in any other disorder.	The much-quoted SAINT remission of about 90% and the IPS guideline's 86.4% are two reported figures for the SAME OPEN-LABEL protocol, not two designs; neither is the randomised 52.5% MADRS reduction at 4 weeks. Attribute each figure to its document; never merge. A non-inferiority result licenses substitution on time, not a claim of greater efficacy.

THE TRAP: THE ENTRY COUNTS IN COLUMN TWO ARE IN DIFFERENT CURRENCIES AND DO NOT CONVERT

Three selective serotonin reuptake inhibitors is a count of ONE drug class for ONE indication. Four or more adequate antidepressant treatments counts treatments of any class, and adequate qualifies the response as well as the treatment. Two or more adequate pharmacologic trials is a third unit, and RANZCP fixes no number while also requiring a failed psychological treatment. An unattributed threshold is right only by accident, so name the body in the same breath as the number.

ACTIVE-TREATMENT NETWORK COMPARISONS, AND WHERE NO DIRECT COMPARISON EXISTS

A NETWORK analysis reviewed by the IPS rTMS guideline, which pools direct with indirect evidence and is never the result of a direct trial, found ECT in one specified form of delivery associated with higher response than high-frequency left rTMS, than continuous theta-burst stimulation and than deep TMS; that form belongs to the ECT and neurostimulation notes.
No trial read for this chapter compares tDCS directly against rTMS, so the only comparison between those two is the position NICE gave each, one step apart.
Efficacy and acceptability are separate axes, and this frame carries the efficacy one alone; an answer that reports only it is describing half the choice a patient makes.

DECLARED GAP, NOT FILLED: THE INDIAN POSITION FOR FIVE OF THE SIX MODALITIES

rTMS: no CDSCO classification entry, no Indian product or import licence for a named system, and no Indian indication-specific authorisation appears in any source read here, so a foreign clearance list must not be read across and the chapter does not assert that rTMS is approved in India for anything.
tDCS: transcranial electrical stimulation systems are classified by CDSCO at Neurological serial 127, Class B, which is classification and not approval, and no Asian country except Singapore and South Korea has approved such devices. No Indian clinical practice guideline sets indications or a place in the ladder for tDCS.
DBS: the deep brain electrical stimulation system is listed by CDSCO at Neurological serial 33, Class D, again a classification recording no approval or indication for use. Whether any Indian centre implants for a psychiatric indication, any product licence, any authorisation, and the cost in rupees, are all undocumented.
MST is not documented as available in India, and the Indian position on accelerated iTBS and imaging-targeted protocols is undocumented in the same way: no centre report, no Indian trial, no cost, and no statement on whether the imaging-targeting step is feasible in routine Indian services.

WHAT THIS PLATE DELIBERATELY DOES NOT CARRY. The procedural act for every modality named here belongs to the ECT and neurostimulation notes and is absent from this frame by design: electrode placement, stimulus dosing, motor-threshold referencing, coil specification and the peri-procedural apparatus. This frame is utility alone.

COLOUR KEY

- teal = a stated guideline or regulatory position, and the strength it carries
- grey = no position stated in the sources read here, and the boundary note
- coral = where the position is negative or null, and the trap that follows
- dashed outline = a declared absence in the sources read here

Fig 37.9 Neurostimulation utility across disorders: six modalities placed by entry point, by indication, and by the strength of the statement that grants each one. The entry column is a set of thresholds and not a sequence, because the counts are in different currencies and do not convert: no failure count at all for ECT where a rapid or definitive response is needed to avert harm, three limbs bound together in NICE TA59, two or more adequate pharmacologic trials for rTMS on a weak recommendation, three selective serotonin reuptake inhibitors for deep brain stimulation in obsessive-compulsive disorder, and four or more adequate antidepressant treatments for vagus nerve stimulation. Each modality is then graded disorder by disorder with its negative rungs printed beside its positive ones: rTMS is null in panic disorder and in clozapine-refractory schizophrenia, both named tDCS depression trials are null, and the VA/DoD guideline is against deep brain stimulation and vagus nerve stimulation outside research. ECT is the only modality in the frame with a randomised superiority result against drug treatment in depression. The Indian position is a declared absence for five of the six, and the procedural act for all six belongs to the ECT and neurostimulation notes. (IPS ECT and rTMS guidelines 2023; NICE TA59, HTG396, HTG382, HTG551; VA/DoD 2022; RANZCP 2020; UK ECT Review Group; CREST-MST; CORRECT-BD; THREE-D; Stanford Neuromodulation Therapy; US FDA HDE H050003 and PMA P970003/S050; CDSCO.)

A course is a treatment episode, not a treatment plan, and the aftercare decision is examined more often than it is taught. Where depression has responded to a course, the current NICE depression guidance directs that antidepressant medication or a psychological intervention be started or continued to prevent relapse and provide ongoing care, and that lithium augmentation of antidepressant medication be considered. Read the second limb precisely: the augmentation offered for consideration is of the antidepressant, and the modal is permissive, not directive.

GOOD TO REMEMBER

The commonest omission on a real ward is not in the decision to treat, it is in what is written on the discharge summary the week the course ends. A response achieved and then left unprotected is the predictable relapse, and the guidance names two protective options and one augmentation to consider, rather than leaving the choice to habit.

18.6 Cognitive adverse effects, and the gap in how they are measured

This is the half of the answer that separates a pass from a good mark, because the cognitive evidence is genuinely two-sided and most candidates give only one side. The Indian guideline names **retrograde amnesia** as one of the most distressing adverse effects and, in the same breath, as one that is difficult to measure using objective cognitive tests, and it places that statement in its guidance on monitoring during a course rather than in its guidance on consent.

The largest pooled analysis of objective cognitive testing found a course of treatment followed by recovery and then improvement past baseline, on twenty-four cognitive variables:

- From 0 to 3 days after treatment, a significant decrease in 72 percent of the variables, with effect sizes from -1.10 to -0.21.
- Between 4 and 15 days, all but one confidence interval included zero.
- Beyond 15 days, no negative effect sizes at all, and 57 percent of the variables significantly improved, with effect sizes from 0.35 to 0.75.

That is reassuring and it is incomplete, and the incompleteness is stated by the authors themselves: no standardised retrograde amnesia test was identified anywhere in the pooled studies. The recovery finding therefore does not cover the domain the Indian guideline calls the most distressing.

PITFALL

Answering a patient who reports lost memories of their own life by citing the objective recovery data. The measured domains recovered; the domain the patient is describing was not among the standardised measures. Offering the reassurance anyway tells a patient their experience contradicts the evidence when in fact the evidence is silent on it, and it is the fastest way to lose a person's trust in a treatment that has just worked.

What does exist for that domain is a separate meta-analysis of **autobiographical memory** measured against controls, pooling nine studies, 432 patients and 173 controls in adults with unipolar or bipolar moderate-to-severe depression, with maintenance courses excluded.

Immediately after treatment the patients had greater autobiographical memory loss than the controls, a moderate effect, with an interval from 0.35 to 0.75. From the patients' own side, a systematic review of patient perspectives found reported persistent memory loss ranging from **29 percent to 55 percent** across studies, and, unlike reported benefit, that rate did not appear to depend on whether the study was clinician-based or patient-based. The convergence matters: this is not a complaint that disappears when you ask a clinician instead.

18.7 Cardiac and systemic risk, stated at the right size

The systemic risk of the modality is small, real, and almost always misquoted in one of two directions. The largest meta-analysis, pooling 82 studies, 106,569 patients and 786,995 treatments, found the three most commonly reported **major adverse cardiac events** to be acute heart failure, arrhythmia and acute pulmonary edema. The pooled incidences were 24 per 1,000 patients for acute heart failure, **25.83 per 1,000 patients for arrhythmia** and 4.92 per 1,000 patients for acute pulmonary edema. Restated against treatments rather than patients they fall to 2.44, 4.66 and 1.50 per 1,000 ECT treatments respectively, and the difference between those two denominators is the whole point: a patient receiving a course is exposed repeatedly, so the per-patient figure is the one that answers "what is my risk", and the per-treatment figure is the one that answers "what is the risk of this session".

One further caution about that literature, and it is a correctness point rather than a nuance. The pooled mortality construct in the same review is all-cause death within 30 days of treatment, which is not the same thing as death attributed to the treatment. Writing it as a mortality rate for the modality overstates it, and an examiner who knows the source will read that as a failure to read it.

18.8 Regulation of the modality in India, and what is actually practised

Indian law regulates this modality directly, which is unusual and therefore heavily examined. **Unmodified electroconvulsive therapy** means the treatment given without muscle relaxants and anaesthesia, and Section 95(1) of the Mental Healthcare Act 2017 prohibits it for any person with mental illness, and prohibits the treatment in minors. The prohibition in minors is not absolute: Section 95(2) provides that notwithstanding sub-section (1), if the psychiatrist in charge of a minor's treatment is of the opinion that the treatment is required, it shall be done with the informed consent of the guardian and prior permission of the concerned Board. Separately, Section 94(3) of the same Act provides that nothing in its emergency-treatment section permits a medical officer or psychiatrist to use the treatment as a form of emergency treatment, while emergency treatment itself is limited to 72 hours, extendable to seven days in a declared disaster or emergency.

IN INDIA

Hold the two rules together and a real clinical decision falls out of them. The Indian guideline places the treatment first-line precisely where a rapid response is needed to avert harm, which is exactly the presentation that arrives out of hours and feels like an emergency. The statute closes the emergency-treatment route for this modality. So the lawful path to treating that patient quickly is the consent or supported-admission pathway, started at once rather than after the emergency window has been used up on something else. The clinician who plans for the consent route on the first evening treats sooner than the one who reaches for the emergency provision and discovers it is unavailable.

Regional practice diverges from the northern-hemisphere guidance in a way that is worth stating plainly in an Indian answer. Across South Asia the treatment is widely practised, most commonly delivered with brief-pulse devices in its modified form; schizophrenia, not affective illness, is the commonest indication in the region; electroencephalographic monitoring of seizures is rarely practised; and practice guidelines for the modality are available only in India.

TRAP **Quoting the British position on schizophrenia in an Indian answer without naming the divergence.** One body declines to recommend general use in schizophrenia on the state of the evidence; in this region schizophrenia is the commonest indication for which the treatment is given, and the Indian guideline's indication table is written accordingly. Both are true statements about different things, and the marks go to the candidate who states both and labels which is which, not to the one who picks the tidier of the two.

An answer on this modality is complete only when it carries the trigger, the evidence at its true strength in each disorder, the memory cost including the domain nobody has measured well, and the Indian statute that governs how it may lawfully be given.

19 · Repetitive transcranial magnetic stimulation (rTMS): mechanism, protocols and indications

A ten-mark answer on rTMS is an answer about utility, not about apparatus. Repetitive transcranial magnetic stimulation is the neuromodulation technique most likely to reach an ordinary outpatient, and the marks fall on four questions: what the magnetic field does to cortex, where the modality sits once drugs have failed, which disorders it has earned an indication in and on what strength of evidence, and what happens to the patient once the course finishes. The delivery technique itself, and the machinery of safe administration that surrounds it, are set out in the ECT and neurostimulation notes; what follows is the modality and its utility across disorders.

The commonest way to lose these marks is to answer as though rTMS were one treatment with one indication. It is a family of protocols whose evidence is graded disorder by disorder, at strengths that run from strong to frankly negative, and the bodies that recommend it do not agree on how much treatment should have failed before it is offered. Separate the disorders, attribute the guidelines, and the answer organises itself.

19.1 What the field does, and the sentence that defines the modality

In **transcranial magnetic stimulation** the stimulator is discharged to produce a strong time-varying magnetic field, briefly peaking at **1 to 2.5 Tesla** and lasting 1 ms at most. A field of that strength crosses scalp and skull with little attenuation, and that is the property that makes the technique non-invasive. Non-invasive is not the same as unfelt: scalp discomfort, scalp pain and headache are recognised effects of the stimulation itself, and they belong in the consent conversation rather than in a claim that the treatment cannot be felt. Where the field meets conducting tissue it induces an electric current, and where that induced current is large enough it depolarises cortical neurons directly rather than through a peripheral route.

A single pulse probes the motor system and is a neurophysiological measurement tool. What makes the technique therapeutic is repetition: pulses delivered repeatedly within a session, and sessions repeated over weeks, shift cortical excitability in a direction that outlasts the stimulation itself. That outlasting shift, not the individual pulse, is what the treatment is betting on, and it is the reason a course is counted in sessions rather than in minutes of stimulation. The frequency of delivery influences the direction of that shift, which is why the evidence below is reported as high-frequency or low-frequency stimulation to a named cortical target rather than as rTMS in the abstract. Frequency does not settle the direction on its own: intensity, the pattern in which the pulses are grouped, the amount of stimulation delivered, the target and the state of the cortex at the time all bear on which way the after-effect goes, and the parameters themselves belong to the procedural notes. Treat the pairing of high frequency with excitation and low frequency with inhibition as a heuristic that organises the literature, not as a law.

RULE **Therapeutic rTMS is deliberately sub-convulsive.** The seizure is the event the treatment is designed to avoid, not the event it is designed to produce. That single sentence separates rTMS cleanly from convulsive therapy and from magnetic seizure therapy, whose purpose is the opposite, and every safety statement made about rTMS follows from it.

19.2 Where the modality sits once drugs have failed

The Indian Psychiatric Society issues its own clinical practice guideline for the therapeutic use of rTMS in neuropsychiatric disorders, and its governing statement about placement is not about any single disorder. rTMS is to be used as an **adjunct** to other conventional treatments, and the guideline records that this is how it had been given in almost every meta-analysis it reviewed, across almost every condition it examined. The point is methodological before it is clinical: most of the effect sizes quoted for rTMS are effect sizes for rTMS added to something else, and reading them as the performance of rTMS given alone overstates them. Check the individual estimate before treating it as evidence for augmentation.

RULE **The evidence prices rTMS as an addition, not a substitution.** In treatment-resistant depression specifically, the same guideline finds response to rTMS better when it is added as an **augmentation** to antidepressants than when it is given stand-alone. The prescription that continues alongside the course is part of the treatment being tested.

PITFALL

Stopping or tapering the antidepressant when the rTMS course begins, on the reasoning that the patient is now getting something better. The comparison behind the guideline's own position was augmentation against stand-alone treatment, and it favoured augmentation. A patient whose drug is withdrawn at the first session is no longer receiving the treatment the evidence describes.

Where the bodies genuinely differ is on how much failure should precede the referral, and the differences are not cosmetic. Read the table below across the third column as much as the second: a threshold carries the authority of the statement that sets it, and two of these four are explicit that their statement is weak.

BODY	WHERE IT PLACES RTMS	STRENGTH OF THAT STATEMENT
Indian Psychiatric Society rTMS guideline	An adjunct to conventional treatment across the conditions reviewed, with the depression evidence graded strong	Guideline recommendation resting on graded evidence
VA/DoD major depressive disorder guideline	Suggests offering rTMS to patients showing partial or no response to **two or more adequate pharmacologic treatment trials**	Weak recommendation, very low certainty
RANZCP mood disorders guideline	Before rTMS is considered, failure to respond to a reasonable number of adequate trials of both pharmacotherapy and psychological treatment	Graded CBR, consensus-based rather than evidence-based
NICE	rTMS for depression may be used with normal arrangements for clinical governance and audit	Short-term efficacy adequate, clinical response variable

TRAP **Quoting a single line of therapy for rTMS as though the field agreed on one.** One body counts failed drug trials and fixes a number. Another requires failed drug and failed psychological treatment but deliberately fixes no number, and grades that requirement as consensus rather than evidence. The regulatory clearance that first put the device into practice was written against a different bar again from either. An unattributed threshold is therefore right only by accident: state whose threshold you are using in the same breath as the number.

19.3 Depression: the indication the modality rests on

The Indian guideline grades the evidence for rTMS in depression as strong, with pooled effect sizes for improvement in depression severity spanning **0.302 to 0.83**. It also gives the two dichotomous outcomes an examiner is most likely to ask for: pooled odds ratios for response between 3.26 and 3.64, and for remission between 2.45 and 4.63. Quote both kinds of number if the question allows it, because they answer different questions. The effect size says how far the depression score moves; the odds ratio compares the odds of crossing a response or a remission threshold under active stimulation against the odds under sham. It is not a statement of how

many times more likely this patient is to respond, because an absolute probability needs the sham event rate, and no such rate is reported here. A wide effect-size range with a consistent odds ratio is exactly what a heterogeneous but genuinely active treatment looks like.

United Kingdom guidance arrives at a compatible destination by a different route, and its wording is worth memorising rather than paraphrasing. NICE judges the evidence on short-term efficacy adequate although the clinical response is variable, and states that rTMS for depression may be used with normal arrangements for clinical governance and audit. Both halves matter. The permissive modal is not an endorsement of first-line use, and the phrase about normal arrangements is a positioning statement rather than an efficacy claim: it says the procedure needs no special oversight, not that it works better than the alternatives. Note also that the depression guideline itself sets no rTMS recommendation of its own. NICE NG222 cross-refers to the interventional-procedures guidance IPG542, so a candidate who cites the depression guideline as the source of an rTMS recommendation is in fact citing a signpost.

19.4 What the acute course actually buys

A course buys a response; it does not settle the illness. Among patients who respond to rTMS for depression at the end of the acute course, the proportion still holding that **sustained response** falls steadily across the following year.

66.5%	52.9%	46.3%
THIRD-MONTH RESPONSE HELD	SIXTH-MONTH RESPONSE HELD	TWELFTH-MONTH RESPONSE HELD

Read those three figures as one statement rather than as three facts. Roughly two thirds of initial responders are still responders a quarter of a year later, and rather under half of them at a year, with confidence intervals that widen as the follow-up lengthens because fewer studies report the later points. Two consequences follow for how the treatment is offered. First, the acute course is a planning point rather than a discharge point, and the maintenance question has to be raised before the last session rather than after the relapse. Second, these three figures measure the durability of response among patients who responded acutely, and response and remission are distinct outcomes, so no claim about durable remission can rest on them. No figure for sustained remission is reported here at all, and the honest answer says so rather than reusing these proportions for it.

> **GOOD TO REMEMBER**
>
> Book the review before the course ends, not after it. The attrition of response is steepest in the first half of the year, so a patient discharged to routine follow-up at three months is being reviewed on the shoulder of the curve rather than on its plateau. Telling the patient the shape of that curve at the outset also protects the alliance if the response fades.

19.5 The indication ladder beyond unipolar depression

The remaining indications are graded individually and the grades are not interchangeable. The examinable skill is to reproduce the ladder with each rung labelled by its own strength, including

the rungs where the evidence is negative, since an answer that lists only the positive indications is describing a marketing brochure rather than a guideline.

CONDITION	EVIDENCE POSITION	WHAT THE NUMBERS AND THE COURSE LOOK LIKE
Acute bipolar depression	Moderate and positive	Effect size 0.302, odds ratio for response 2.72
Obsessive-compulsive disorder	Moderately positive	Pooled Y-BOCS SMDs 0.501 to 0.79, across 10 to 30 sessions given over 1 to 12 weeks
Generalised anxiety disorder	Significant effect	Pooled effect sizes 1.45 to 1.87; depression accompanying the anxiety improves too, at SMD 1.65
Panic disorder	Not found effective	No effect demonstrated on the pooled analysis
Treatment-resistant auditory hallucinations	Moderate and positive	Pooled SMDs 0.24 to 0.51, across 4 to 40 sessions delivered up to 8 weeks
Negative symptoms of schizophrenia	Moderate to large and positive	SMD 0.49 to 0.64; high-frequency stimulation to the left DLPFC and more than 10 sessions superior to sham
Clozapine-refractory schizophrenia	No significant effect	Three randomised trials of rTMS augmentation, not significant for positive symptoms, negative symptoms or cognition
Tobacco use disorder	Positive	High-frequency rTMS targeted at the left DLPFC reduces both craving and the amount consumed
Other substance use disorders	Not consistent	The evidence does not converge across substances

Two rows on that ladder repay closer reading. The schizophrenia entries are not one finding: negative symptoms carry the strongest signal in the disorder, resistant hallucinations a smaller but real one, and **clozapine-refractory schizophrenia** carries no significant effect at all on any of the three outcome domains examined. That last row is the one candidates forget, and it is precisely the clinical situation in which a desperate team is most tempted to add a device. Treatment resistance in schizophrenia and the place of clozapine within it are handled as a topic of their own; what belongs here is only that adding rTMS on top of clozapine did not separate from sham.

The obsessive-compulsive row carries a disagreement that should be stated rather than resolved. The Indian guideline's own reading is that the evidence for deep TMS in obsessive-compulsive disorder is inconclusive and insufficient, while the regulatory record reproduced inside that same

document lists a **deep TMS** clearance for obsessive-compulsive disorder in 2018. These are different currencies rather than contradictory facts: a regulator asks whether a device is reasonably safe and effective for a labelled use, a guideline panel asks whether pooled trial data reach a grade. Print both, attributed to their sources, and say which question each is answering.

TRAP **Writing that rTMS works in anxiety disorders.** It has a significant effect in **generalised anxiety disorder**, with the depression that accompanies it improving as well, but it was not found effective in panic disorder. Candidates generalise across a diagnostic block because the ladder is memorised as a list of positives, and the negative rungs are where an examiner separates a memorised list from an understood one.

19.6 Where the modality stops, and how it compares with the alternatives

The Indian guideline draws one explicit boundary: it does not recommend rTMS for acutely suicidal patients. That is a positioning statement about tempo rather than about efficacy in depression, and it is the single most examinable sentence in the Indian document, because it identifies the presentation in which this modality is not the answer however strong its depression evidence is elsewhere.

In a network meta-analysis reviewed by the same guideline, in the comparisons between two active treatments, ECT was associated with higher response than high-frequency left rTMS, than continuous theta-burst stimulation, and than deep TMS. A network estimate pools direct and indirect evidence, so report it as a network comparison and never as the result of a trial that compared the two treatments directly. The analysis specified one particular form of ECT delivery, and that specification belongs to the ECT and neurostimulation notes rather than here. A separate analysis in the same reviewed set reported bilateral rTMS more acceptable than ECT, and, in treatment-resistant depression, more effective than deep brain stimulation. Efficacy and acceptability are separate axes; the two findings sit beside each other rather than against each other, and an answer that reports only the first is describing half the choice a patient is actually making.

IN INDIA

The controlling document for this modality is the Indian Psychiatric Society's own rTMS guideline, and it settles two decisions a candidate should be able to state without hesitation. First, the Indian guideline places rTMS as an adjunct to conventional treatment rather than a substitute for it, and that is this guideline's position rather than a universal restriction, since rTMS is not confined to augmentation everywhere. Keep the concomitant antidepressant or antipsychotic stable through the course where that is clinically feasible; if it has to change, the safety of continuing is reviewed rather than assumed, and the parameter side of that review belongs to the procedural notes. Second, it explicitly does not recommend rTMS in acute suicidality, so the acutely suicidal presentation is the boundary at which this modality stops and the alternatives are weighed on their own evidence rather than on an extension of this one.

19.7 Safety, stated at the right size

Safety belongs inside a utility answer rather than appended to it, because the whole clinical case for rTMS is that it delivers a moderate effect at a very low cost in risk. Seizure is the serious adverse event, and its measured rate is very low: a large safety survey reported **7 seizures per 100,000 sessions** overall, from 24 seizures in over 300,000 sessions, rising to 33 per 100,000 sessions among people carrying an identified risk factor. The consensus that reports those figures is careful about them: the survey data cannot be considered more than approximate and may carry reporting biases, so quote the order of magnitude with the caveat attached rather than the decimal with false precision.

The controlled evidence in depression separates the serious from the merely common, and that separation is the shape of the answer.

- Across 53 sham-controlled trials in depression there was no increased risk of serious adverse events.
- There was no increased risk of drop-outs attributable to an adverse event, so tolerability did not differ from sham on the measure that matters for completing a course.
- There was a significantly greater risk of mild, transient, non-serious events, which is where the real difference from sham sits and what consent should describe.
- The same guideline reports no signal of affective switch with rTMS protocols used to treat depression, which is the question a bipolar-spectrum patient's team will ask first.

19.8 The regulatory picture, and rTMS in Indian practice

United States clearance began in depression and widened by indication over the following sixteen years: 2008 for the first TMS system in major depressive disorder, 2013 for the first deep TMS system in the same indication, 2018 for deep TMS in obsessive-compulsive disorder and separately for the first theta-burst-capable system in major depressive disorder, 2020 for smoking cessation, 2021 for deep TMS in major depressive disorder with anxiety, 2022 for an accelerated protocol system in major depressive disorder, and 2024 for major depressive disorder in adolescents aged 15 to 21. Theta-burst and the accelerated protocols are substantial topics in their own right; here they matter only as the shape of the widening, which has moved from one diagnosis, to new diagnoses, to new populations.

None of that is an Indian authorisation, and the distinction has to be made explicitly rather than left to the reader. No device classification, product or import licence, or indication-specific Indian authorisation for any named rTMS system is stated in this account, and a foreign clearance list must not be read across as one. The defensible examination sentence is that rTMS is in established clinical use in India while the Indian regulatory position is not something to assert from another country's labelling.

What is documented is the service picture. TMS services began to be offered by private practitioners in India from the mid-2000s, initially concentrated in metropolitan cities, and treatment centres are now established in smaller cities as well; the concerns that persist are variability in standards of care, high treatment costs, and limited accessibility for the general

population. Training infrastructure has grown alongside the services: a one-year fellowship in **non-invasive brain stimulation** has been offered since 2016, aimed at developing clinical and research expertise in the area. For an Indian answer, that pairing is the honest one to give, because it names both the spread of the modality and the unevenness of the standards under which it has spread.

MNEMONIC

Disorder, Threshold, Adjunct, Durability

The four things a modality answer on rTMS must settle: which disorder, and therefore which grade of evidence applies; whose threshold of prior failure is being quoted; that the treatment is added to conventional care rather than substituted for it; and what proportion of responders are still responding a year later.

The closing position is best stated without inflation, and a recent Indian review, Praharaj in the Indian Journal of Psychiatry in 2025, supplies it whole rather than by halves: the exact role of TMS in the treatment of psychiatric disorders remains uncertain, although it appears to offer benefit in mild to moderate cases, particularly for the short-term treatment of depression. Print that as the review's own position, with its author and year attached, and then set the guideline evidence beside it rather than in place of it, because the two are grading different things. The review is grading severity of illness. The guidelines are grading prior treatment failure, and prior failure is not a severity band: a patient can fail two adequate trials of moderate illness, and a patient with severe illness can be drug-naive. The Indian guideline grades acute depression strong and finds augmentation superior to stand-alone rTMS in treatment-resistant depression, and the VA/DoD statement is written for patients with partial or no response to two or more adequate pharmacologic trials. None of that says anything about how severe the illness was, so none of it refutes the review's severity clause. Hold both, each attributed to the body stating it, and note that the grade everywhere else runs disorder by disorder. That is not a weak answer. It is the accurate one, and it is what lets a candidate hold the strong depression grading and the negative clozapine-refractory row in the same paragraph without contradiction.

rTMS is an adjunct with its strongest evidence in depression, a disorder-by-disorder grading everywhere else, no recommendation in acute suicidality, and a response that has to be maintained rather than assumed.

20 · Transcranial direct current stimulation (tDCS)

This is the modality that is easiest to describe and hardest to place. Transcranial direct current stimulation passes a weak, steady current between contacts on the scalp. It is simple to deliver and it has been trialled across most psychiatric diagnoses, and that combination is precisely why a full-length answer on tDCS is not an answer about promise. The marks sit in four places, and a candidate who supplies only the first loses most of them: what the technique is, in a family that has an explicit boundary; what it does to a neuron, which is not what the seizure-

based and magnetic techniques do; what the best-designed depression trials returned when they asked the question properly; and where guidance and regulators, having read those trials, have actually put it.

The procedural technique is set out in the ECT and neurostimulation notes and is not repeated here. What follows is the modality itself and the honest evidence position across disorders, which is a different examination question and is answered with different material.

20.1 What the technique is, and the family it belongs to

tDCS belongs to a defined family, **low-intensity transcranial electrical stimulation**, and the family is bounded by an explicit dose envelope rather than by an indication or a mechanism. That is an unusual way to define a treatment and it is worth pausing on: the boundary is a set of physical limits on the current delivered, so anything inside those limits belongs to the family and anything outside it is a different technique with a different safety literature. The safety and guidance paper for low-intensity transcranial electrical stimulation sets a ceiling on total current, a frequency span, and a ceiling on how long a single session may run.

4 mA or less	0 to 10,000 Hz	up to 60 minutes per session
TOTAL INTENSITY CEILING	FREQUENCY SPAN OF THE FAMILY	LONGEST SINGLE SESSION

The frequency span is the limb that explains the family's shape. Direct current sits at the zero-frequency end of it, which is what the middle word of the technique's name means: the current does not oscillate, it holds a constant direction for the whole session. The alternating-current members of the same family occupy the rest of the span, and for those the intensity ceiling is expressed peak-to-peak rather than as a steady value, because an oscillating waveform swings either side of zero and its amplitude has to be quoted across the full swing. One envelope therefore holds several techniques that look different at the bedside and are grouped within one consensus safety framework for low-intensity transcranial electrical stimulation.

Two cautions about the envelope. It is a definitional and safety boundary agreed by the field, not a legal classification that holds in every jurisdiction, and it is not a therapeutic prescription: nothing in it says what current a given patient should receive, and how a session is actually configured belongs to the procedural chapter named above. And a boundary drawn around a family is not evidence about any member of it. The envelope tells you what tDCS is; it tells you nothing at all about whether it works.

20.2 The mechanism, and why it is the answer's spine

tDCS acts by **subthreshold neuronal polarisation**, which is to say it does not itself generate action potentials. The current shifts the **membrane potential** of neurons in the field without directly bringing them to threshold, and thereby changes how likely they are to fire in response to synaptic input they are already receiving. The safety and guidance paper states this as the mechanistic contrast with rTMS and with ECT, and it is the single most examinable sentence in the topic.

RULE **tDCS does not directly trigger action potentials. It shifts membrane potential and changes how likely neurons are to fire in response to synaptic input they are already receiving.** Every other property of the modality follows from this. A technique that biases ongoing activity rather than driving it produces no directly evoked motor response to titrate against, unlike a suprathreshold technique, and its effects may be state-dependent, varying with what the cortex is already doing while the current is applied. Write this sentence first and the rest of the answer organises itself around it.

Two corollaries follow from the definition rather than from any separate finding, and stating them as derivations is safer than presenting them as results. First, the effect can be expected to be state-dependent: if the current only changes a probability of firing, then what the neuron is already doing is part of the outcome, which is why the same delivered current need not be the same intervention in two different patients or in the same patient on two different occasions. Second, the distance between this mechanism and a clinical response is very long. A membrane-potential shift is several inferential steps away from a change in mood or in a symptom score, and that gap is where the modality's evidence problem lives.

20.3 Polarity, and why the classroom rule is not a law

Every candidate has met the rule that the anode excites and the cathode inhibits. It is a useful first approximation and it is how the technique is usually introduced. It is also the single most confidently over-stated fact in this topic. The same passage that establishes the subthreshold mechanism qualifies it directly: the mechanisms of low-intensity transcranial electrical stimulation vary with the properties of the current applied and with how that current is applied, so the polarity heuristic cannot be printed as a universal rule.

TRAP **Writing "anodal stimulation is excitatory, cathodal is inhibitory" as a law of the modality.** Candidates make this error because the heuristic is how the technique is taught, because it is easy to remember, and because it has the satisfying symmetry of a real physical law. The source that establishes the mechanism explicitly qualifies the direction of effect as varying with the properties of the current and with the way it is delivered. The safe examination form is to state the heuristic as a heuristic, say what it is derived from, and then state the qualification in the same breath. An examiner who has read the current guidance is looking for the qualification, not for the rule.

20.4 Depression: the two named trials

Depression is where the modality has been tested hardest, and two trials are named here because each was designed to answer a different question and each returned a result the candidate must be able to state precisely. In **ELECT-TDCS**, a **single-centre trial**, tDCS did not show non-inferiority to escitalopram over **10 weeks** in the treatment of depression, and it was associated with more adverse events than the drug comparator. That limit belongs to the result itself and should travel with it whenever the trial is cited.

The second is DepressionDC. A published analysis from that trial records that the main trial did not show beneficial antidepressant effects of active tDCS over **sham** tDCS. The provenance of that sentence is worth stating rather than hiding: it is drawn from a subsequent peer-reviewed

report by the same investigators and not from the primary publication, so it is a faithful account of the main trial's result as its own authors describe it.

AXIS ON WHICH THE TWO DESIGNS DIFFER	ELECT-TDCS	DEPRESSIONDC
What active tDCS was compared against	An active drug comparator, escitalopram	Sham stimulation
Question the design is built to answer	Non-inferiority: is the device not meaningfully worse than the drug	Superiority: does active stimulation beat sham
What the null answer licenses	That the trial could not exclude a meaningful shortfall against the drug; not that the technique is inert	That active stimulation did not beat sham in that trial's population; not that it must fail in every population
Limit the source states alongside the result	A single-centre trial	The statement used here comes from a later report by the same investigators

20.5 What those results license, and what they do not

A failed trial and an ineffective treatment are not the same object. This is the reasoning the neuromodulation questions are actually testing, and it is where most answers go wrong in the same direction. A **non-inferiority** design asks whether the new intervention falls short of the comparator by more than a pre-specified margin. Failing to demonstrate non-inferiority means the trial could not rule that shortfall out. It does not mean the intervention did nothing, it does not mean the comparator was superior, and it certainly does not mean the modality has been refuted. A sham-controlled superiority design asks a different question, and a null there is a statement about that contrast, in that population, at the parameters that trial used.

The discipline this imposes on an answer is specific. State each trial by name with its own comparator and its own result, and stop. Do not aggregate two negative results into a verdict on the modality, and do not soften them into a vague statement that the trials have been disappointing, which sounds cautious but is an unsourced summary claim in disguise. Nothing here supports a statement about what the tDCS depression literature as a whole has shown, and an answer that offers one has left the evidence behind.

TRAP **Collapsing two named negative trials into "tDCS does not work".** This is a different error from the polarity one and it is made by the better-read candidates, who know the trials and over-read them. Two results, one against an active drug and one against sham, are two data points with their own designs, populations and stated limits. Neither was designed to answer whether the modality has utility in any other disorder, and neither can be extended to one. The correct answer names the trials, states what each showed, and says plainly what remains unsettled.

20.6 Tolerability, and the assumption a low current invites

The intensity envelope makes tDCS feel like a low-risk intervention, and the temptation is to write that it is well tolerated and move on. The one head-to-head comparison available here does

not support that shortcut. In the trial against escitalopram, tDCS was associated with more adverse events than the drug, and the events reported in that comparison were specific rather than trivial.

- Higher rates of skin redness in the tDCS group.
- Higher rates of tinnitus.
- Higher rates of nervousness.
- **New-onset mania** in the tDCS group.

> **PITFALL**
>
> Reading a low current as a low adverse-event rate. The two are separate questions and this modality separates them: the envelope is at the low end of what electrical stimulation can deliver, and the one trial read here that compared it directly against a drug still recorded more adverse events on the device arm. The clinical error that follows is describing tDCS to a patient as the option with fewer side effects, which is a comparative claim the evidence read here contradicts rather than supports.

The appearance of new-onset mania in the reported events is the item to carry away, because it is the one that changes what is asked before treatment and requires monitoring for emergent hypomanic or manic symptoms during treatment. A mood-disorder history is part of the picture for any intervention given with antidepressant intent, and the reported event list, not the current intensity, is the honest basis for that conversation.

20.7 Where guidance places the modality

Guidance positions are examinable in their own right, and this one has a distinctive shape. NICE places tDCS for depression under **special arrangements**, and the arrangements it names are clinical governance, consent, and audit or research. That is a category in its own right and it should be read as one: it is neither a prohibition nor an endorsement of routine use, but a position that permits the procedure while requiring that its use be governed, consented and either audited or conducted as research. NICE puts tDCS for depression in a lower position than it gives rTMS, and the comparison is part of the answer, because a candidate who states the tDCS position without that relative placement has given a fact without its meaning.

That relative placement is the only direct cross-modality comparison available here, and it must not be treated as the sole clinical basis for choosing when a patient could plausibly be offered either non-convulsive technique. There is no head-to-head trial here putting tDCS against rTMS, so the choice between them is not made on a direct comparison of outcomes. What this account holds is the position each has been given by the body that reviewed them, and on that basis tDCS is the one carrying the additional governance, consent and audit conditions. Other considerations that bear on the choice, including each technique's own adverse-effect profile, its contraindications and what the service can actually deliver, sit outside what the evidence read here covers. Saying so plainly is more defensible than manufacturing a comparison the evidence does not contain.

GOOD TO REMEMBER

If tDCS is being offered for depression, the consent conversation is not the same conversation as for an established treatment. The position it rests on names consent and audit or research as conditions of use, which means the patient is being told that the intervention sits in a governed category rather than in routine practice, and that the service is expected to record what happens. Building that into the conversation is not extra caution; it is the position itself.

20.8 The Indian position: classified, not approved

India has a regulatory record for this technology and it is easy to misread. Transcranial electrical stimulation systems are classified by the Central Drugs Standard Control Organisation as a **medical device**, at **Neurological serial 127**, Class B. The classification entry's own intended-use text names transcranial direct current stimulation and transcranial alternating current stimulation for psychiatric or neurological therapy, so the technology is recognised by name. Classification is not approval. It places a device in a risk class for regulatory purposes; it does not certify that the device works for a named indication. The wider position is that no Asian country other than Singapore and South Korea has approved tES devices, so India holds classification without indication approval.

IN INDIA

There is no Indian national clinical practice guideline setting indications, evidence grades or a place in the treatment ladder for tDCS in psychiatric disorders. That absence is the operative fact, not an omission from this account. The consequence for practice is concrete: the position an Indian clinician is working from is a guidance position derived from NICE plus the named trial evidence, and it should be presented to the patient and recorded in the notes as exactly that. Because classification at Class B is not approval for a psychiatric indication, the decision to offer tDCS is a governed, consented and documented one taken on the trial record, not one taken on the strength of an approved indication or of a domestic guideline that does not exist.

MNEMONIC

Envelope, Polarity, Trials, Position

The four beats of a full tDCS answer, in the order that earns marks. Envelope: the family and the physical limits that define it. Polarity: the subthreshold mechanism and the heuristic that must not be stated as a law. Trials: the named depression trials with their own comparators and their own limits. Position: what guidance permits and what Indian regulation has and has not done. An answer missing the last two beats is a description, not an answer.

Set beside the other techniques, tDCS occupies a distinctive place: a mechanism that is well characterised and modest, an evidence base in depression that is specific, named and unfavourable in the two trials read here, a guidance position that permits use only under governance, consent and audit or research, and an Indian regulatory record that recognises the device without approving an indication for it. Each of those four is a fact a candidate can state

and defend. The temptation is to replace them with a general impression of a promising low-cost treatment, and that impression is the one thing in this topic that nothing here supports.

> tDCS biases firing rather than driving it, and every honest claim about it - mechanism, evidence and regulatory position - is smaller and more specific than the enthusiasm the technique attracts.

21 · Deep brain stimulation and vagus nerve stimulation in psychiatry

Two of the neurostimulation answers require an implanted system. Deep brain stimulation and vagus nerve stimulation sit at the far end of the treatment sequence, reached only after documented pharmacological failure, and a full-mark answer on either is marked on four things: what the device is, where it acts, which disorders it has been used in, and how strong the evidence actually is. Two errors cost most of the marks. The first is promoting a regulatory status into an efficacy claim, which is easy to do because the regulatory language is written to sound like endorsement. The second is answering with the operation instead of the modality. The procedural technique, the consent and legal frame around the operation itself, and the ethics and cost-equity arguments are set out in the ECT and neurostimulation notes; what follows is the modality and its utility across disorders.

21.1 A circuit therapy, not a reversible lesion

The first conceptual correction is the one examiners look for, and it is worth opening an answer with. **Deep brain stimulation** is now better described as a **circuit therapy** and not as a reversible lesion, though it was initially conceptualised as one, a framing that made it the ablative operations' safer successor and that still shapes how referrers and families talk about it. The delivered pulses trigger action potentials in axons both orthodromically and antidromically, and the new patterns of synaptic activity this generates overwrite the pathological ones.

The distinction is not semantic, and it changes three things in an answer. It explains why the effect can be modulated or withdrawn entirely without any tissue having been removed, which is what separates this modality from the ablative procedures it descended from. It explains why the unit of treatment is a network rather than a structure, so that the clinical question is which circuit is being driven and not which nucleus is being destroyed. And it explains why the same anatomical target can support claims in more than one disorder: if the mechanism is the imposition of new activity patterns on a circuit, then any disorder in which that circuit is implicated becomes a candidate indication, which is exactly how the psychiatric target list was assembled.

RULE **A device authorisation names who may be treated, not how well the treatment works.** Both of the psychiatric authorisations discussed here are defined by a population and by a count of prior failed drug treatments. Neither is a statement about effect size, and the strength of the underlying evidence has to be read separately from the wording of the licence. Read the pathway, then read the trials, and never let the first stand in for the second.

21.2 The psychiatric targets, and why obsessive-compulsive disorder leads

Target selection is where a modality answer earns its structure, and obsessive-compulsive disorder is the psychiatric condition in which the target work has gone furthest. In deep brain stimulation for obsessive-compulsive disorder, the **ventral capsule/ventral striatum** and the nucleus accumbens have emerged as the most frequently investigated loci, with three further sites receiving substantial attention.

- Most frequently investigated: the ventral capsule/ventral striatum.
- Equally prominent: the nucleus accumbens.
- Substantial attention: the anterior limb of the internal capsule.
- Substantial attention: the bed nucleus of the stria terminalis.
- Substantial attention: the subthalamic nucleus.

Two observations make this list examinable rather than a memory exercise. First, the sites are not independent of one another: they cluster in and around the ventral striatal and capsular territory through which cortico-striato-thalamo-cortical traffic passes, which is the anatomical reason a single disorder produced several candidate targets rather than one. Second, the fact that five targets remain under investigation is itself the evidence position. A field that had settled its target would not still be studying five, and an answer that names one target as the target has misread the state of the literature. Say which are studied most, say the rest are actively studied, and resist the temptation to declare a winner.

21.3 The regulatory footing in obsessive-compulsive disorder

The United States authorisation for this indication is a **Humanitarian Device Exemption** and not a **premarket approval**, and the difference is the whole point. US FDA granted it as HDE H050003, with a decision date of 10 February 2009, for bilateral stimulation of the anterior limb of the internal capsule, as an adjunct to medications and as an alternative to anterior capsulotomy, for chronic severe treatment-resistant obsessive-compulsive disorder in adults who have failed **at least three** selective serotonin reuptake inhibitors.

Every clause in that sentence is load-bearing. The indication is bilateral. It is adjunctive, so the device is added to medication rather than replacing it. It is positioned clinically as an alternative to anterior capsulotomy; this wording does not establish that capsulotomy was a trial comparator. It is restricted to adults. And it carries the crispest prior-failure threshold anywhere in this modality group: a specific count, of a specific drug class, before the device becomes an option at all.

TRAP **Writing that deep brain stimulation is approved for obsessive-compulsive disorder.** A humanitarian device exemption is a humanitarian pathway carrying a probable-benefit standard, not a demonstration of effectiveness, and it is expressly not a premarket approval. Candidates make this error because the everyday word for any regulatory clearance is approval, and because the indication text reads exactly like a licence. Name the pathway, then the standard it applies, then the threshold it sets. That single sentence separates a candidate who has read the authorisation from one who has heard about it.

21.4 Deep brain stimulation in depression: the controlled evidence

Depression is the second psychiatric indication for which this modality has been pursued seriously, and it is the one on which the controlled evidence is clearest and least encouraging. The randomised, double-blind BROADEN trial did not show significant differences between active and sham stimulation at six months, and it was subsequently halted because of a low probability of success on an interim **futility analysis**.

Both halves of that sentence teach. A trial that fails to separate from sham has not shown that the device does nothing; it has shown that this design, in this population, at this time point, could not demonstrate that it does something, and open-label improvement in the same patients is entirely compatible with that result. But a trial stopped for futility is a stronger negative than a trial that merely misses significance, because the interim analysis has already concluded that continuing to the planned endpoint was unlikely to change the answer. An answer that reports the sham comparison without the futility stop understates the result; an answer that reports the futility stop as proof of inefficacy overstates it.

The guideline consequence follows directly. VA/DoD recommends against using deep brain stimulation for major depressive disorder outside a research setting, and does so at strong strength. Note the verb and the strength together, because both are the substance of the recommendation: this is a recommendation against, not a recommendation to consider carefully, and it is issued at the guideline's higher tier rather than its weaker one.

21.5 Vagus nerve stimulation: what the indication actually says

The second implanted modality is licensed the other way round. Where the obsessive-compulsive disorder authorisation came through a humanitarian route, **vagus nerve stimulation** in depression holds a premarket approval: US FDA P970003 supplement S050, with a decision date of 15 July 2005. The device is indicated for the adjunctive long-term treatment of chronic or recurrent depression, in patients 18 years of age or older who are in a major depressive episode and have not had an adequate response to **four or more** adequate antidepressant treatments.

Three words in that indication do most of the examinable work. Adjunctive means the device is used in addition to other treatment rather than in place of it, and the two authorisations say it with different force: the obsessive-compulsive disorder exemption reads adjunct to medications, while this one reads only adjunctive long-term treatment. Long-term is unusual in a device indication and is not decoration: it distinguishes the intended treatment horizon from an acute intervention, and it should not be relabelled as maintenance treatment, since the same indication requires that the patient be in a major depressive episode. Chronic or recurrent restricts the population by illness course, not merely by severity, so a first episode meets the illness-course limb only if it is chronic.

PITFALL

Counting prior failures loosely, or interchangeably between the two devices. The obsessive-compulsive disorder pathway counts a specific class, selective serotonin reuptake inhibitors. The depression pathway counts, in the indication's own words, four or more adequate antidepressant treatments, and adds the word adequate twice, to the treatments and to the response; that indication text does not define those treatments by class. A clinician who tallies four failed drugs of mixed classes for the obsessive-compulsive disorder pathway, or three selective serotonin reuptake inhibitors for the depression pathway, has counted the wrong thing against the wrong threshold and has not established candidacy for either.

21.6 The evidence behind the vagus nerve stimulation indication

This is the single most misremembered fact in the whole modality group, and it is worth stating plainly: the acute sham-controlled trial behind the depression indication was negative. In the D-02 acute phase, **15.3 percent (17 of 111)** of the active stimulation group were considered responders against 10.0 percent (11 of 110) of the sham group, a difference that was not statistically significant, with p equal to 0.238. Support for the indication came instead from long-term open treatment and from nonrandomised comparative evidence, with the limitations that implies; in that uncontrolled long-term phase the reported 12-month HAM-D response was 29.8 percent and remission 17.1 percent.

Set those figures side by side and the structural problem in the evidence is visible without any statistical apparatus. The controlled comparison, which is the only part of the programme in which sham stimulation is present, produced a small and non-significant separation. The larger response proportion comes from a phase with no comparator, in which the passage of time, continued pharmacological treatment, regression to the mean and expectation all remain in the estimate. A candidate who quotes only the long-term figure has quoted the uncontrolled number as though it were a treatment effect.

The most recent evidence has not resolved this. RECOVER, described as the largest neuromodulation trial in psychiatry to date, did not meet its primary endpoint on preliminary reports, although clinically meaningful improvements were reported in depressive severity, functioning and quality of life. That combination, a missed primary endpoint alongside reported secondary improvements, is the recurring shape of this literature and should be described as such rather than resolved in either direction.

TRAP **Assuming a premarket approval implies a positive pivotal sham-controlled trial.** Here it does not. The acute controlled comparison was negative and the support came instead from long-term open treatment and nonrandomised comparative evidence, and the largest trial since has not met its primary endpoint either. This is a genuinely different error from mistaking the humanitarian pathway for approval: there the pathway was misread, here the pathway is read correctly and the evidence behind it is assumed rather than checked.

21.7 How the guideline bodies read the two devices

Guideline positions are more consistent across bodies than the target literature is, and they are the safest thing to reproduce under time pressure. A systematic review of clinical practice guidelines on deep brain stimulation for obsessive-compulsive disorder found that eight guidelines recommended it be used only after all other treatment options had failed to alleviate symptoms, and that one guideline did not recommend it beyond a research setting. In depression the positions are more restrictive still: VA/DoD recommends against deep brain stimulation outside research at strong strength, and suggests against vagus nerve stimulation outside research at weak strength. NICE places implanted vagus nerve stimulation for treatment-resistant depression under **special arrangements** for clinical governance, consent, and audit or research, in guidance now numbered HTG551 and formerly IPG679, published 12 August 2020.

at least 3 SSRIS FAILED, US OCD PATHWAY	**4 or more** ANTIDEPRESSANT TRIALS, US VNS INDICATION	**8** GUIDELINES: ONLY AFTER ALL ELSE FAILS	**1** GUIDELINE: RESEARCH SETTING ONLY

Read the modal verbs against each other and a hierarchy appears. Recommends against is stronger than suggests against. Special arrangements is neither a prohibition nor an endorsement: it permits use while requiring the governance, consent and audit machinery that ordinary practice does not, which is what a body does when it is unwilling to close a door and unwilling to open it. Compared on the axes an examiner is likely to supply, the two devices separate cleanly.

AXIS	DEEP BRAIN STIMULATION	VAGUS NERVE STIMULATION
Lead psychiatric indication in the regulatory record	Chronic severe treatment-resistant obsessive-compulsive disorder in adults, adjunctive	Chronic or recurrent depression, adjunctive and long-term
Regulatory pathway	Humanitarian device exemption, not a premarket approval	Premarket approval supplement
Prior-failure threshold named in the authorisation	Failed at least three selective serotonin reuptake inhibitors	No adequate response to four or more adequate antidepressant treatments
Controlled evidence in depression	BROADEN did not separate from sham at six months and was stopped for futility	Acute sham-controlled phase negative; the largest trial to date did not meet its primary endpoint
Position in major depressive disorder	Recommends against outside research, strong	Suggests against outside research, weak

GOOD TO REMEMBER

When a patient or family raises a device after several failed trials, two answerable questions come before any discussion of the operation. How many adequate prior treatments have there actually been, counted against the threshold that belongs to that indication. And is the device being sought as a replacement for medication, which neither authorisation supports: the obsessive-compulsive disorder exemption is written as an adjunct to medications, and the depression approval as adjunctive long-term treatment.

21.8 The position in India, and the question the Act leaves open

The Indian regulatory picture is thinner than candidates expect, and the honest answer says so rather than filling the space. What is documented is a risk classification. CDSCO lists the deep brain electrical stimulation system in the Neurological category of its classification of medical devices as Class D. That is all it is. The cited document is a classification list; it records no approval and no indication for use, so it cannot be quoted as evidence that the device is authorised here for any psychiatric condition, and it says nothing at all about vagus nerve stimulation in depression.

What is not documented, and therefore not stated here, is whether any Indian centre currently implants these devices for a psychiatric indication and for how many patients, whether any product licence exists for a named system, whether any indication-specific authorisation exists, and what the treatment costs in rupees. Those are real questions and an answer is stronger for naming them as unanswered than for supplying a plausible figure. The equity and cost arguments around access to implanted neuromodulation are developed in the ECT and neurostimulation notes.

IN INDIA

Before any implanted device is offered for a psychiatric indication, establish whether the Board approval requirement applies. Under section 96 of the Mental Healthcare Act, 2017, **psychosurgery** is restricted: it requires the informed consent of the person on whom it is performed and, in addition, approval from the concerned Board. The Act does not define psychosurgery, so whether a reversible implanted system for refractory obsessive-compulsive disorder falls inside that restriction cannot be settled from the Act itself. The consequence of the ambiguity is concrete and binary. Either Board approval must be obtained before proceeding, or it need not, and no clinician should resolve that question informally. State the question, state what turns on it, and do not answer it.

MNEMONIC

Site, Stage, Sham, Stance

Four retrieval slots that fit any device answer, including one you have not revised. Site: what is stimulated and on what mechanistic account. Stage: what must have failed, and how many times, before the device is on the table. Sham: what the controlled comparison actually showed, as distinct from the open-label follow-up. Stance: what the guideline bodies recommend, in their own modal verbs and at their own strengths.

Each of these authorisations is defined by a count of prior failed drug treatments, and in depression the sham-controlled evidence for both devices did not separate.

22 · Magnetic seizure therapy and other emerging neurostimulation modalities

Every examination cycle asks for a modality nobody in the room has used. The emerging neurostimulation question is not a request for enthusiasm. It asks whether a candidate can say what a device does, which physical variable it manipulates, which disorders it has been tested in, and how strong that evidence actually is, and can do all four without inflating an investigational technique into an established treatment. Magnetic seizure therapy anchors this group because it has been tested against an active comparator in randomised trials, and because the shape of that evidence is itself examinable. The remaining modalities are best learned by what each one varies: the waveform of the current, the form of energy delivered, or the timing of stimulation. The procedural technique for every device named here, and the apparatus that surrounds it, is set out in the ECT and neurostimulation notes.

22.1 Magnetic seizure therapy, and the inversion that defines it

Magnetic seizure therapy uses the physics of transcranial magnetic stimulation. A rapidly changing magnetic field crosses scalp and skull without the attenuation an externally applied electrical current suffers, and induces a current in the cortex beneath. Repetitive transcranial magnetic stimulation, which has its own section, uses that induced current subconvulsively: the entire discipline of the technique lies in modulating cortical excitability while staying below the point at which a seizure is provoked. Magnetic seizure therapy takes the same class of hardware and reverses the intention. It is delivered so as to produce a generalised seizure, and the seizure is the therapeutic act, exactly as it is in electroconvulsive therapy.

RULE **Magnetic seizure therapy is magnetic stimulation hardware carrying convulsive intent.** One sentence, and it is the whole definition. Repetitive transcranial magnetic stimulation is delivered so as not to provoke a seizure; magnetic seizure therapy is delivered so as to provoke one. Every other feature of the modality, from where it is performed to what it is compared against, is downstream of that inversion.

Two consequences follow at once, and both are answerable without any parameter. First, magnetic seizure therapy is a convulsive treatment and therefore needs the peri-procedural support that a convulsive treatment requires, delivered in an appropriately equipped treatment

and recovery setting. That requirement places it beside electroconvulsive therapy in service terms and separates it from the subconvulsive magnetic techniques it shares hardware with, however similar the two look on a photograph. Second, for a trial testing non-inferiority and the cognitive trade-off against the established convulsive treatment, the correct comparator is electroconvulsive therapy rather than sham or a subconvulsive protocol, because the claim being tested there is not that seizures help but that this way of producing them costs the patient less. A trial asking a different question, such as whether magnetic seizure therapy has a specific effect beyond the procedure itself and expectancy, would call for a credible sham instead; the comparator follows the question. The settings that achieve the inversion, and the whole procedural apparatus around them, belong to the ECT and neurostimulation notes and are not repeated here.

22.2 Why the seizure is meant to be different, and what is actually known

The rationale is physical. An electrical current applied at the scalp must cross tissue of high resistance before it reaches the brain, and the path it takes is determined as much by that tissue as by where the operator puts it. A magnetically induced current is not attenuated in the same way, so in principle the initiating volume of cortex can be smaller and more predictable. If the seizure begins in a smaller volume, the argument runs, it may engage less of the brain, and the cognitive cost may be correspondingly lower. That hypothesis is the reason the modality exists.

What the evidence supports is narrower than the hypothesis, and the difference between the two is where marks are won. The IFCN safety consensus on transcranial magnetic stimulation records that seizures induced by magnetic seizure therapy are less robust and differ in their physiological expression from those induced by electroconvulsive therapy, and it records that on animal evidence only, not on human evidence. That is a statement about how the seizure behaves, not a statement about how far it spreads through the human brain. Human electrophysiological or imaging evidence on whether the magnetically induced seizure propagates less widely, and on how much of any observed cognitive advantage a propagation difference would explain, is not established.

TRAP **Writing focality as the established explanation for the memory difference.** It is the mechanism candidates reach for because it is tidy and because it was the design rationale. The honest position is that the animal data show the seizures differ, that the human trials show a difference on the autobiographical memory outcome they measured, and that the causal link between the two has not been demonstrated in people. An answer that states the hypothesis as hypothesis and then names the missing evidence reads as better trained than one that asserts the mechanism.

22.3 The confirmatory randomised evidence

The design is as examinable as the result. A **non-inferiority** trial sets aside the question of whether the new treatment is better. It asks whether the new treatment falls short by no more than a pre-specified margin, so that its advantages elsewhere can legitimately decide the choice. That question is the right one here, because nobody claims magnetic seizure therapy produces

more remission than electroconvulsive therapy; the claim is that it costs the patient less memory for a comparable chance of getting well.

The CREST-MST trial put that claim to a confirmatory randomised test against ultrabrief electroconvulsive therapy, on 239 randomised participants.

27.8%	22.5%	17.3%	2.7%
REMISSION, ULTRABRIEF ECT ARM	REMISSION, MST ARM	MEMORY WORSENING, ECT ARM	MEMORY WORSENING, MST ARM

Read the two pairs in opposite directions. On remission the gap was **5.3%** in favour of electroconvulsive therapy, and non-inferiority of magnetic seizure therapy was nonetheless established, with p=0.048 and a 95% CI -4.4 to 14.9. On worsening of **autobiographical memory** the difference ran the other way and was of a different order of magnitude, at p=0.0003. So the trial delivers precisely the trade the modality was designed to offer, and it delivers it with the remission point estimate favouring the older treatment.

TRAP **Reading non-inferiority as equivalence, or worse, as superiority.** Non-inferiority is a bounded conclusion: it says the true difference is unlikely to be worse than the margin the trial declared before it started. It does not say the two treatments produced the same remission rate, and in this trial they did not. A candidate who writes that magnetic seizure therapy matched electroconvulsive therapy on remission has misread the design; the defensible sentence names the direction of the difference and then says non-inferiority was established.

22.4 The bipolar pilot, and what a pilot can and cannot settle

The CORRECT-BD pilot asked the same question in bipolar depression and found the same pattern. Remission was reached by 6/20 (30%) in the electroconvulsive therapy group and 5/25 (20%) in the magnetic seizure therapy group, while clinically important worsening of autobiographical memory occurred in 6/27 (22.2%) against 2/28 (7.1%). Remission favoured the older treatment numerically; memory favoured the newer one.

Two features of those figures deserve attention because they are what an examiner is testing. The study randomised n=55. It was not powered for efficacy, so the remission comparison is descriptive and no inference about relative efficacy can be drawn from it. And the denominators are not the same across the two outcomes, because the participants contributing a remission judgement and those contributing a memory assessment are not the same set. Quoting a percentage from a pilot without saying that it came from a pilot is the commonest way this material is got wrong, and it is a failure of scientific reading rather than of memory.

22.5 Reading the two trials as one answer

An answer that lists both trials and stops has done half the work. The two are not two attempts at the same question; they occupy different rungs of the evidence ladder and they license different sentences. Setting them out on shared axes makes that explicit.

AXIS	CONFIRMATORY NON-INFERIORITY TRIAL	RANDOMISED PILOT IN BIPOLAR DEPRESSION
Question it was built to answer	Whether magnetic seizure therapy falls short of the active comparator by no more than the declared margin	Whether the same pattern appears in bipolar depression, and whether a definitive trial is feasible
Comparator	Ultrabrief electroconvulsive therapy	Ultrabrief electroconvulsive therapy
Direction on remission	Numerically favoured electroconvulsive therapy	Numerically favoured electroconvulsive therapy
Direction on autobiographical memory	Favoured magnetic seizure therapy	Favoured magnetic seizure therapy
What the result licenses	A conclusion of non-inferiority on remission, bounded by the confidence interval	Hypothesis generation and planning of a definitive trial
What it does not license	Any claim that magnetic seizure therapy produces more remission	Any conclusion about relative efficacy at all

GOOD TO REMEMBER

When a patient or a family asks about the memory-sparing alternative to electroconvulsive therapy, the honest account carries both halves of the result and not just the attractive one. CREST-MST found less worsening on its specified autobiographical memory measure, and the memory figures set out above are that result; it should not be widened into a claim about cognition as a whole. The chance of remission was numerically lower for magnetic seizure therapy in the confirmatory trial and in the pilot alike, and in the confirmatory trial the conclusion drawn was non-inferiority, not equality. Presenting the memory figure alone converts a defensible clinical discussion into a promise the evidence does not support.

22.6 Beyond the seizure: the transcranial electrical waveforms

The rest of the group varies something other than the intention. The transcranial electrical modalities are distinguished here by the waveform of the current applied at the scalp; their intensity, montage, electrode geometry and the resulting distribution of current in the head may also differ. Guidance on low intensity transcranial electric stimulation endorsed by ESBS and IFCN separates the family by waveform in this way.

- Transcranial direct current stimulation, which has its own section, applies a current of constant intensity and polarity.
- Transcranial alternating current stimulation, **tACS**, alternates between polarities at a steady sinusoidal frequency.
- Transcranial random noise stimulation, **tRNS**, alternates at a randomly changing frequency.

The distinction is not cosmetic, and the reasoning behind it is worth carrying. A constant current shifts the resting membrane potential of the underlying cortex in one direction and so biases the

likelihood of firing. An oscillating current can, in principle, interact with the brain's own rhythms, which makes the frequency chosen a hypothesis about which rhythm is disordered rather than an arbitrary setting. A randomly varying current makes no such frequency-specific claim; its rationale is the addition of stochastic input to a network. Three different theories of how to help a sick circuit therefore sit inside one apparently uniform family, and an answer that says so has understood the field.

What cannot be stated here is how well any of them works in any named disorder. No disorder-level efficacy figure for alternating current or random noise stimulation is available to this section, and the correct examination position is to describe the mechanism and the family structure, name the modality as investigational, and decline to attach an effect size to it.

22.7 Transcranial focused ultrasound

Focused ultrasound is the only technique in this group that changes the form of energy rather than the shape of a current. Acoustic energy can be steered so that it converges on a small volume at depth while passing through the tissue in front of it at much lower intensity, which is the property that makes it interesting: it offers a non-invasive route to targets that surface stimulation cannot reach, and it does so without an implanted device.

The point that must be got right is that the same physical technology spans two entirely different clinical acts, separated by intensity. **HIFU** used for ablation typically employs acoustic intensities **greater than 200 W/cm2**, whereas **low-intensity focused ultrasound** used for neuromodulation generally operates **below 100 W/cm2**. One destroys tissue at the focus and the change it produces is permanent; the other is intended to modulate the activity of tissue it leaves intact. A candidate who writes about focused ultrasound without naming that split has written about two procedures as though they were one, and the resulting answer cannot be corrected by adding detail, because the error is categorical.

22.8 Adaptive and closed-loop stimulation

The last variable available to an emerging modality is not where or how hard, but when. Conventional deep brain stimulation, which has its own section, runs open-loop: the device delivers what it was programmed to deliver, and the programme changes only when a clinician changes it. **Adaptive, closed-loop deep brain stimulation** senses a physiological signal from the patient and varies stimulation in response to it, so the device is titrating itself against the state of the circuit rather than against a clinic appointment. Conceptually this is the most significant shift in the group, because it moves the treatment from a fixed prescription to a feedback system.

Its regulatory position is precise and is the part most often stated wrongly. The US FDA granted adaptive deep brain stimulation as an optional programming feature on an existing deep brain stimulation system under PMA P960009 supplement S478, with a decision date recorded as 20 February 2025, and that authorisation is for Parkinson's disease. There is no corresponding authorisation for any psychiatric indication. The adaptive principle is being studied in psychiatry, but the approval that exists sits outside it.

PITFALL

Telling a patient with severe refractory illness, or a family that has read about it, that closed-loop brain stimulation is now approved and available for their condition. The approval is for Parkinson's disease. Describing a movement-disorder authorisation as though it extends to psychiatric indications raises an expectation that cannot be met and is a factual error in the record, not merely an optimistic phrasing.

22.9 Answering the emerging-modality question honestly

The group has no shared mechanism, so it needs a shared method of answering. For any device in it, the same three questions produce a structured reply, and they can be applied to a modality a candidate has never seen.

- Which physical variable does it change, and what does the change intend to do to the circuit?
- What has it been compared against, in whom, and how many were studied?
- What does the resulting evidence license as a sentence, and what does it not?

MNEMONIC

Intent, Waveform, Energy, Timing

The four axes on which the emerging modalities separate. Magnetic seizure therapy changes the intent of magnetic stimulation from subconvulsive to convulsive. Alternating current and random noise stimulation change the waveform. Focused ultrasound changes the form of energy. Adaptive stimulation changes the timing. Nothing in the list asserts an efficacy claim, which is exactly why it is safe to reproduce under examination conditions.

Applied to this material, that method gives an answer with different strengths in different places, and the differences are the substance. Magnetic seizure therapy carries randomised comparative evidence against an active treatment, in a confirmatory trial and a pilot that point the same way on remission and on memory. The electrical waveform variants and focused ultrasound neuromodulation are described here by mechanism and by physical parameter, and no disorder-level efficacy figure is attached to them, because none is available to this section. Adaptive stimulation has a regulatory position, and that position is outside psychiatry.

One further limit belongs in the answer rather than in a footnote. No regulatory classification, product licence or indication authorisation for a magnetic seizure therapy device in India could be established from the sources behind this section, nor any report of an Indian centre delivering it clinically or in research; the same absence holds for therapeutic alternating current and random noise stimulation and for focused ultrasound neuromodulation. The position that is defensible in an Indian examination and in an Indian clinic is therefore that magnetic seizure therapy is investigational and is not documented as available here, stated as an absence of evidence rather than as an assertion of unavailability.

Each emerging modality is defined by the single variable it changes, and its evidence is worth exactly what its comparator and its sample size make it worth.

23 · Theta-burst stimulation and accelerated TMS protocols (SAINT) as recent neurostimulation advances

Theta-burst is not a new machine, it is a new rhythm. Three separate innovations sit behind the phrase "recent advances in transcranial magnetic stimulation", and the popular account of the field runs them together: a change in how the pulses are patterned, a change in how the sessions are spaced across the calendar, and a change in how the cortical target is chosen. Theta-burst stimulation is the first. Accelerated protocols are the second. Imaging-based targeting is the third. The best-known protocol in this area is all three at once, which is why its headline result is quoted more often than it is understood. An answer that keeps the three axes apart, and attaches every published number to the design that produced it, will outscore one that recites remission rates.

The procedural side of all of this, including how the instrument is specified and how treatment intensity is referenced before a course begins, belongs to the ECT and neurostimulation notes and is not repeated here. What follows is the modality answer: what theta-burst is, where the accelerated schedules came from, what the trials actually demonstrated, where the two guideline positions currently sit, and what remains undocumented.

23.1 What theta-burst stimulation is

Theta-burst stimulation is a patterned form of repetitive transcranial magnetic stimulation rather than a modality in its own right: instead of delivering single pulses at one even frequency, the stimulator delivers short bursts of high-frequency pulses and repeats those bursts at the **theta rhythm**. The parameters the majority of the literature has used are those originally described by Huang and colleagues, 50 Hz bursts of 3 pulses repeated at 5 Hz. The rationale is borrowed rather than invented. Burst patterning at theta frequency is the arrangement that was used in animal preparations to induce lasting change in synaptic efficacy, and the argument for importing it into human cortical stimulation is that a patterned train induces plasticity more efficiently than an evenly spaced one of comparable size. That argument is a rationale, not a demonstrated human mechanism, and the clinical case for the pattern has had to be made separately in trials.

The consequence a candidate should be able to state in one line is administrative before it is biological. Because the pulses arrive in dense bursts, a treatment session runs to roughly **3 minutes**, dramatically shorter than the conventional session it substitutes for. The parent modality and the clinical setup around it are otherwise similar, but the dose structure, the session duration and the tolerability are not identical, which is precisely why the substitution had to be tested rather than assumed.

RULE **Theta-burst changes the pattern, not the modality.** The same instrument delivers the same kind of stimulation to the same patient; what changes is how the pulses are arranged in time, and with that arrangement the dose structure, the length of the session and the tolerability of it. It follows that nothing about the parent technique is superseded, and that every therapeutic claim specific to the pattern needs its own evidence. Two forms of the pattern are used, an intermittent one and a continuous one, and the depression evidence that follows is entirely about the intermittent form.

23.2 THREE-D, and what a non-inferiority result licenses

The trial that moved intermittent theta-burst out of the physiology laboratory and into routine depression practice was **THREE-D**. In patients with **treatment-resistant depression**, intermittent theta-burst was **non-inferior** to 10 Hz repetitive transcranial magnetic stimulation, with low dropout numbers and similar side-effect, safety and tolerability profiles in both arms. That sentence, and the design word inside it, is the examinable core.

The quantities matter because the argument is statistical before it is clinical. HRSD-17 scores improved from 23.5 to 13.4 in the high-frequency arm and from 23.6 to 13.4 in the theta-burst arm, an adjusted difference 0.103 with a lower 95% CI of -1.16 and p=0.0011, tested against a pre-specified non-inferiority margin set at 2.25 points. Read what that structure does: the confidence bound sits well inside the margin, so the trial excludes a loss of effect larger than the margin allows. It does not measure a gain.

Tolerability was close but not identical, and the asymmetry is worth carrying. Dropouts across the two arms were 6% and 8%, headache was the commonest adverse event in both at 64% and 65%, and pain during treatment was rated higher with theta-burst, 3.8 vs 3.4 out of a possible 10, p=0.011. A shorter session is not automatically a more comfortable one, and a patient who declines the technique on those grounds is not being unreasonable.

TRAP **Writing that this trial showed theta-burst to be better, or that it showed the two to be identical.** A non-inferiority design with a pre-specified margin can only exclude a loss larger than that margin, and the primary analysis here was per-protocol rather than intention-to-treat. The result licenses substituting the shorter protocol for the longer one on grounds of time and throughput. It does not license a claim of greater efficacy, and an absent difference is not the same finding as demonstrated equality.

23.3 The regulatory footing came before the recommendation

Theta-burst acquired a device footing before it acquired a graded guideline recommendation, and knowing the sequence explains why the technique is widely available while the formal appraisals remain unsettled. An Indian review of the evolution of transcranial magnetic stimulation services in psychiatry tabulates the regulatory record for neuropsychiatric indications, and its 2018 entry is the MagVita TMS Therapy System with TBS (Tonica Elektronik, MagVenture) for major depressive disorder, recorded as the first theta-burst entry in that table.

A regulatory entry of this kind is a statement about one named system and its labelled indication. It is not a graded appraisal of a body of evidence, and it says nothing about accelerated scheduling or about imaging-based targeting, which came later and are separately evidenced.

Candidates conflate the two because both are commonly described as approval; keeping them apart is a cheap mark.

23.4 Accelerated protocols and the Stanford protocol

"Accelerated" is a statement about the calendar, not about the pulse. An **accelerated protocol** delivers more than one session in a day so that a course which would otherwise occupy weeks is compressed into days. Three things can therefore be varied independently, and a protocol description that does not specify all three is incomplete:

- how many sessions are given in a day;
- how many days the whole course runs;
- how the cortical target is chosen for the individual patient.

The best-developed example is **Stanford Neuromodulation Therapy**, the accelerated **fMRI-targeted** intermittent theta-burst protocol previously called SAINT. Beyond the compressed schedule, its distinctive feature is the target: the region of left DLPFC most functionally anticorrelated with the **subgenual anterior cingulate cortex**, identified in the individual patient by functional imaging rather than assigned by a group rule.

In a double-blind randomised trial reported at a planned interim analysis, the protocol beat sham. The mean percent reduction in MADRS score 4 weeks after treatment was 52.5% with active treatment against 11.1% with sham, in 29 treated (14 active, 15 sham) from 32 enrolled. That denominator is the part which falls out of retelling. The report's own conclusion is that further trials are needed to determine the protocol's durability and to compare it with other treatments, so presenting it as an established treatment misstates its authors' position, not merely the evidence.

52.5% at 4 weeks	11.1% at 4 weeks	29
ACTIVE, MADRS REDUCTION	SHAM, MADRS REDUCTION	PATIENTS TREATED

GOOD TO REMEMBER

SAINT and Stanford Neuromodulation Therapy are the same protocol under two names, the newer replacing the older. A literature search or a guideline read under only one of them will miss half the record, because the trial report uses the newer name while the Indian guidance cites the older. When a referrer or a patient asks about "the Stanford protocol", they are asking about the accelerated imaging-targeted schedule, not about theta-burst patterning as such.

23.5 The remission figure, and which study produced it

The number that reaches patients and referrers is a remission rate near ninety percent, and it did not come from the randomised comparison above. The same report records a remission rate of approximately **90%** after 5 days of open-label treatment. That is a different cohort, a different design and a different outcome measure from the sham-controlled reduction in depression score, and the two quantities answer different questions.

A second figure circulates beside it. The Indian Psychiatric Society guideline on the therapeutic use of repetitive transcranial magnetic stimulation records SAINT remission at **86.4%** in treatment-resistant depression, from a high-dose accelerated protocol running 10 daily sessions for 5 days guided by resting-state functional connectivity imaging, and states in the same passage that such protocols remain to be tested in controlled studies. These are two reported figures for the same open-label protocol, appearing in two documents. Print whichever your source supports, attach the source to it, and do not average or reconcile them.

TRAP **Quoting the remission rate as though the randomised trial produced it.** The controlled result is a percentage reduction in a depression rating scale against sham; the remission rate is an open-label figure from treatment without a sham comparator. Candidates merge them because both are large percentages attached to the same protocol name, and the merged sentence reads as though a sham-controlled trial delivered remission in nine patients out of ten. It did not, and the two published open-label figures do not even agree with each other.

PITFALL

Delivering the compressed schedule without the imaging step and calling it the same protocol. Both the trial report and the Indian guidance describe the protocol as functionally imaging-guided, with the target defined in the individual patient. A compressed schedule aimed by a conventional method may be a perfectly reasonable thing to study, but it is not the protocol the published remission figures came from, and those figures should not be offered for it in a consent conversation.

23.6 Where the guideline positions stand, and why they diverge

Two guideline positions can be asked for, and they do not agree. The divergence is instructive rather than embarrassing, and an answer that names it is stronger than one that smooths it.

The VA/DoD guideline for major depressive disorder declines to recommend either for or against theta-burst stimulation, recording insufficient evidence, a strength of neither for nor against, and certainty of very low. Note the shape of that entry. It is a statement about the state of the evidence base, not a finding that the technique is ineffective, and reading it as the latter is a common error.

The Indian Psychiatric Society position on accelerated protocols is cautious for a related but separate reason: a meta-analysis restricted to accelerated protocols over the left dorsolateral prefrontal cortex was not found to be associated with significant response, and the SAINT remission figure is explicitly flagged as untested in controlled studies. The caution is attached to acceleration, not to patterning.

The tension is worth stating plainly. The non-inferiority trial showed intermittent theta-burst to be non-inferior to a high-frequency protocol that the VA/DoD guideline does recommend, and yet that guideline still declines to grade theta-burst itself. Guideline appraisals grade a body of evidence for a named intervention; inheriting a recommendation across a comparison, however sound the comparison, is not how they are constructed.

SOURCE OF THE FIGURE	DESIGN THAT PRODUCED IT	WHAT IS REPORTED	WHAT IT LICENSES
The non-inferiority trial	Randomised, per-protocol primary analysis, treatment-resistant depression	Non-inferior to 10 Hz; adjusted difference 0.103 against a margin set at 2.25 points	Substitution of the shorter protocol, on time and throughput
The Stanford protocol trial	Double-blind sham-controlled, planned interim analysis, 29 treated from 32 enrolled	MADRS reduction 52.5% active against 11.1% sham, 4 weeks after treatment	An efficacy signal in a small sample; durability left open
The same report, open-label cohort	Open-label, no sham comparator	Remission approximately 90% after 5 days of treatment	Hypothesis generation only
Indian Psychiatric Society guidance	Guideline citation of the same open-label protocol	Remission 86.4% in treatment-resistant depression	The figure met in Indian reading, flagged untested in controlled studies
VA/DoD depression guideline	Formal appraisal of theta-burst as an intervention	Insufficient evidence; neither for nor against; very low certainty	No recommendation in either direction

23.7 What predicts response, and what that finding is not

The Indian guideline's depression evidence review reports a meta-analysis restricted to theta-burst studies, which identified three features predicting higher response rates:

- at least **1800 pulses per session**;
- **subthreshold intensity**;
- a treatment duration of two weeks or less.

Three cautions travel with that list, and an answer that prints the predictors without them has overstated the source. First, these are predictors recovered from a meta-analysis and reported inside an evidence narrative; they are not issued as a recommended protocol, and the guideline does not present them as one. Second, they were derived in depression, and the same review does not establish them for obsessive-compulsive disorder, schizophrenia or addiction. Third, the direction of the first predictor is easy to invert under exam pressure: the finding is a floor, more pulses in a session rather than fewer, while the duration predictor runs the other way, a shorter course rather than a longer one. How any of these is actually implemented in the treatment room belongs to the ECT and neurostimulation notes.

The accelerated-protocol meta-analysis sits beside this list and has to be reported beside it rather than folded into it. That analysis, restricted to accelerated protocols over the left dorsolateral prefrontal cortex, found no significant association with response. The theta-burst analysis above reports three features associated with higher response rates among the studies it pooled. Neither analysis compares one schedule directly against the other, so neither licenses a statement about

what acceleration adds or removes relative to a conventionally spaced course, and no direct schedule-comparison analysis was found to settle it.

23.8 Utility beyond depression, and what is not on record

Every figure above belongs to depression, and the honest answer says so. No theta-burst efficacy figure for any other disorder is carried here, and the response predictors are explicitly not established outside depression. The disorder-by-disorder indication set for conventional repetitive transcranial magnetic stimulation is taught with that modality and should be answered from there rather than reconstructed under this heading; what can be said about theta-burst outside depression is that the evidence used here does not extend to it, which is a different statement from saying the technique does not work there.

The Indian footing for the accelerated protocols is undocumented rather than negative, and the distinction is worth making explicitly instead of filling in. Whether any Indian centre currently delivers an accelerated or imaging-targeted theta-burst course in clinical service, whether any Indian trial of accelerated intermittent theta-burst in depression has been conducted, what such a course costs a patient here, and whether the functional-imaging targeting step is feasible in routine services are all questions the sources behind this section do not answer. They would be answered by the national neuromodulation body and by the academic centres that run these services. Until they are, the defensible position in an answer is that this is an international recent advance whose Indian position is not on record.

IN INDIA

Where instrument time rather than referral volume is what limits how many patients finish a course, intermittent theta-burst is the throughput decision. The session is short enough that more patients can be treated on one machine in a working day, and the non-inferiority result means the substitution is not paid for in effectiveness that trial could detect. Note carefully what the Indian Psychiatric Society caution attaches to: it attaches to the accelerated, several-sessions-a-day protocols and to the imaging-targeted variant, and not to theta-burst patterning itself, which sits inside the same guideline's depression evidence review as a studied form of the parent technique. Choosing the pattern and declining the acceleration is a defensible position, but name each authority separately instead of summing them into a joint endorsement: the non-inferiority trial supports intermittent theta-burst as non-inferior to the high-frequency protocol, the Indian Psychiatric Society is cautious specifically about acceleration, and the VA/DoD guideline remains neutral on theta-burst itself, neither for nor against it.

MNEMONIC

Design, Denominator, Endpoint, Durability

Four questions to put to any accelerated-stimulation result before quoting it: was there a sham arm, how many patients were actually treated, which scale and which time point produced the figure, and has the effect been shown to last. The Stanford report answers the first three and leaves the fourth open in its own words.

Theta-burst changes the pattern and acceleration changes the calendar: every headline number in this field belongs to one design, and an answer is only as strong as the design it names beside the figure.

Pharmacogenetic and structural drug-reaction syndromes

24 · HLA-B*1502 and HLA-A*3101 testing and carbamazepine, oxcarbazepine and phenytoin severe cutaneous reactions (SJS and TEN)

This is the one prescription in psychiatry that a laboratory result is allowed to veto before the first tablet. Most pharmacogenetics adjusts a dose; here a single allele decides whether a drug is used at all. The marks sit on five questions, asked in this order: which allele, which drug, which population, when the test is worth doing, and how to read a test whose negative result is close to conclusive while its positive result is mostly a false alarm. A candidate who can only say that carbamazepine causes severe skin reactions in Asian patients has answered none of them.

Two alleles carry the topic: one is narrow in reaction and wide in drug, the other wide in reaction and narrow in drug. Three aromatic agents are implicated, and the classic error is that the substitute a candidate reaches for is one of the three. The wider family of idiosyncratic reactions, including drug reaction with eosinophilia and systemic symptoms, is taken up elsewhere in these notes.

24.1 The reaction the test exists to prevent

Stevens-Johnson syndrome and **toxic epidermal necrolysis** are one process separated by area. CPIC defines them by epidermal detachment as a proportion of body surface area, with a formally named overlap band between them, so a case is classified by extent rather than by morphology. The bands are exclusive at their edges, and detachment of exactly 10 per cent is overlap rather than Stevens-Johnson syndrome; the guideline's own prose writes the lower band loosely enough to assign that edge to both, so the strip below prints the boundary as the standard clinical classification sets it.

less than 10%	10 to 30%	more than 30%
BSA, SJS	BSA, OVERLAP	BSA, TEN

Mortality separates just as sharply as the area does. CPIC gives rates typically below 5% for Stevens-Johnson syndrome and above 30% for toxic epidermal necrolysis, with sepsis the most frequent cause of death, and adds that mortality is also related to age, the drug half-life and how early the drug is discontinued. Two of those three are fixed by the time the rash appears; only the third is in the prescriber's hands. Half-life as a pharmacokinetic principle belongs to the Psychopharmacology I: principles and core drug classes notes.

The population weighting is what makes this an Indian examination topic rather than a Western one. The carbamazepine boxed warning estimates these reactions at 1 to 6 per 10,000 new users

in countries with mainly Caucasian populations, with the risk in some Asian countries about 10 times higher. For oxcarbazepine the label reports that the reporting rate of toxic epidermal necrolysis and Stevens-Johnson syndrome, itself generally accepted to be an underestimate because of underreporting, exceeds background incidence rate estimates by a factor of 3- to 10-fold.

GOOD TO REMEMBER

Because mortality tracks how early the drug is stopped, the instruction that changes the outcome is given at the first prescription, not at the first rash: stop the tablet and attend the same day if any rash appears. It costs nothing and is the one part of this topic that does not depend on a genotype being available.

24.2 Two alleles doing two different jobs

The first discrimination is between allele and phenotype. **HLA-B*15:02** is tied to one reaction and, on the original Han Chinese work, to a very large effect: an odds ratio for carbamazepine-induced Stevens-Johnson syndrome and toxic epidermal necrolysis of 1357, with a 95% confidence interval of 193.4 to 8838.3. The same study showed the association is phenotype-specific rather than a general marker of drug allergy: HLA-B*15:02 is not associated with maculopapular exanthema or hypersensitivity syndrome, whereas HLA-A*31:01 is associated with maculopapular exanthema, at an odds ratio of 17.5 with a 95% confidence interval of 4.6 to 66.5. A later systematic review put the same point in clinical language: carriers of HLA-B*15:02 are at strongly increased risk of carbamazepine-induced Stevens-Johnson syndrome and toxic epidermal necrolysis in populations where the allele is common, but not of carbamazepine-induced hypersensitivity syndrome or maculopapular exanthema.

HLA-A*31:01 is the mirror image. CPIC states that, unlike the other allele, it is associated with a wider range of carbamazepine hypersensitivity reactions, including maculopapular exanthema, drug reaction with eosinophilia and systemic symptoms, and Stevens-Johnson syndrome and toxic epidermal necrolysis, in many different populations. Breadth of phenotype does not buy breadth of drug: the phenytoin guideline records that no association between HLA-A*31:01 and phenytoin-induced Stevens-Johnson syndrome and toxic epidermal necrolysis has presently been found, despite the structural similarity between the two drugs.

RULE **Each allele names a reaction and a drug list, and neither list generalises.** One allele is narrow in phenotype and wide in drug, covering three aromatic agents for a single reaction; the other is wide in phenotype and narrow in drug, covering several reactions for carbamazepine alone. An answer that treats them as two interchangeable risk markers for the same thing has lost the structure of the topic.

24.3 Four numbers, four different questions

An odds ratio in the thousands is a statement about association, not about test performance, and the two are routinely conflated. CPIC reports that **HLA-B*15:02 as a predictive test for carbamazepine-induced Stevens-Johnson syndrome and toxic epidermal necrolysis has estimated sensitivity 98% and specificity 97% in the Han Chinese population**. The

population qualifier is part of the number: no equivalent estimate exists for any Indian population. A provenance point travels with them: CPIC cites the 2006 Han Chinese study for those figures, but that study's abstract itself reports only the odds ratio, not sensitivity or specificity, so attribute them to the guideline rather than to the primary paper.

Sensitivity and specificity describe the test; predictive values describe the patient, and in a rare outcome they diverge violently. In the same population, the **positive predictive value** of the allele for carbamazepine-induced Stevens-Johnson syndrome and toxic epidermal necrolysis is **7.7%** and the **negative predictive value** 100%. The allele is therefore close to necessary and nowhere near sufficient, since the great majority of carriers would never have developed the reaction. That asymmetry justifies the strategy: a test used to exclude does not need a high positive predictive value.

The same asymmetry sharpens across the two drugs. CPIC states that the positive predictive value of HLA-B*15:02 for oxcarbazepine-induced Stevens-Johnson syndrome and toxic epidermal necrolysis is estimated at 0.73%, much lower than the 7.7% for carbamazepine-induced disease, while the negative predictive value for both nears 100% in Southeast Asian populations. Note that the guideline does not speak with one voice on that last figure: its Table 2 footnote states flatly that HLA-B*15:02 has a 100% negative predictive value for carbamazepine-induced Stevens-Johnson syndrome and toxic epidermal necrolysis, and that its use is currently recommended to guide use of carbamazepine and oxcarbazepine only, while the narrative section quoted above stops at the word "nears". Print both as the guideline's own, and read the word "only" in the footnote as governing which drugs the test guides, not as qualifying the figure.

For the second allele the arithmetic is entirely different and is often quoted out of scope. CPIC estimates that the positive predictive value and the **number needed to test** for HLA-A*31:01, to prevent one case of all carbamazepine-induced hypersensitivity reactions combined, are most favourable in European populations, at 43% and 47 respectively. Two qualifiers travel with that pair and neither is optional: it is a European estimate, and it is most influenced by maculopapular exanthema rather than by drug reaction with eosinophilia and systemic symptoms. A positive predictive value of that size buys mostly the prevention of a benign rash.

TRAP **Carrying the carbamazepine negative predictive value across to phenytoin.** It is stated for carbamazepine, and the phenytoin guideline says the opposite in terms: a negative HLA-B*15:02 test does not eliminate the risk of phenytoin-induced Stevens-Johnson syndrome and toxic epidermal necrolysis. Candidates generalise because the allele and the reaction are the same, and the drug is the variable they stop watching. Note too that 7.7% appears twice in this literature with different meanings: the carbamazepine positive predictive value, and the proportion of a Taiwanese cohort testing positive.

24.4 Which population, and the figures that do not agree

Population frequency decides whether testing is worth offering; it never decides how to read a result that has come back. Published figures for the same populations differ by source and, for the second allele, by what is counted, so they are set out below by source and by measure rather than as one consensus number.

SOURCE	ALLELE	POPULATION, AS THAT SOURCE NAMES IT	MEASURE AND FIGURE
CPIC carbamazepine and oxcarbazepine guideline	HLA-B*15:02	East Asian; Oceanian; South and Central Asian	Allele frequency 6.9%; 5.4%; 4.6%, the three highest
CPIC carbamazepine and oxcarbazepine guideline	HLA-B*15:02	Japanese; Korean	Frequency less than 1%; less than 2.5%
CPIC carbamazepine and oxcarbazepine guideline	HLA-B*15:02	African; African American, Middle Eastern, Caucasian, Hispanic and South American	Not observed; less than 1%
CPIC phenytoin guideline	HLA-B*15:02	East Asian and Central or South Asian; European	Allele frequency 1% to over 20%; 0 to 1%
Carbamazepine label	HLA-B*1502	South Asians, including Indians	Intermediate prevalence, averaging **2% to 4%, but higher in some groups**
Oxcarbazepine label	HLA-B*1502	India; Korea	Allele frequency up to about 6%; up to about 2%
CPIC carbamazepine and oxcarbazepine guideline	HLA-A*31:01	Caucasian; Hispanic and South American; Japanese; South Korean; South and Central Asian	Allele frequency 3%; 6%; 8%; 5%; 2%
Carbamazepine label	HLA-A*3101	Japanese, Native American, Southern Indian (for example, Tamil Nadu) and some Arabic ancestry	Expected carriage more than 15%

Three divergences in that table are examinable in their own right. The two regulatory labels disagree with each other about India, one giving an intermediate prevalence of 2% to 4% for South Asians including Indians and the other an allele frequency up to about 6% for India; both are current, so the honest answer names both rather than averaging them. The two CPIC documents describe the same populations with a single aggregate in one and a wide range in the other, so a figure must be attributed to the guideline it came from. Most importantly, the second allele's Indian figures are not comparable: the guideline reports an allele frequency in South and Central Asians, the label an expected carriage in Southern Indian ancestry. Different quantities over different groupings, so these are two measurements and are never converted into one. Note also that the guidelines write the colon in an allele name and the labels omit it; both spellings denote the same allele, and a stem may use either.

24.5 What the instruments instruct, and at what strength

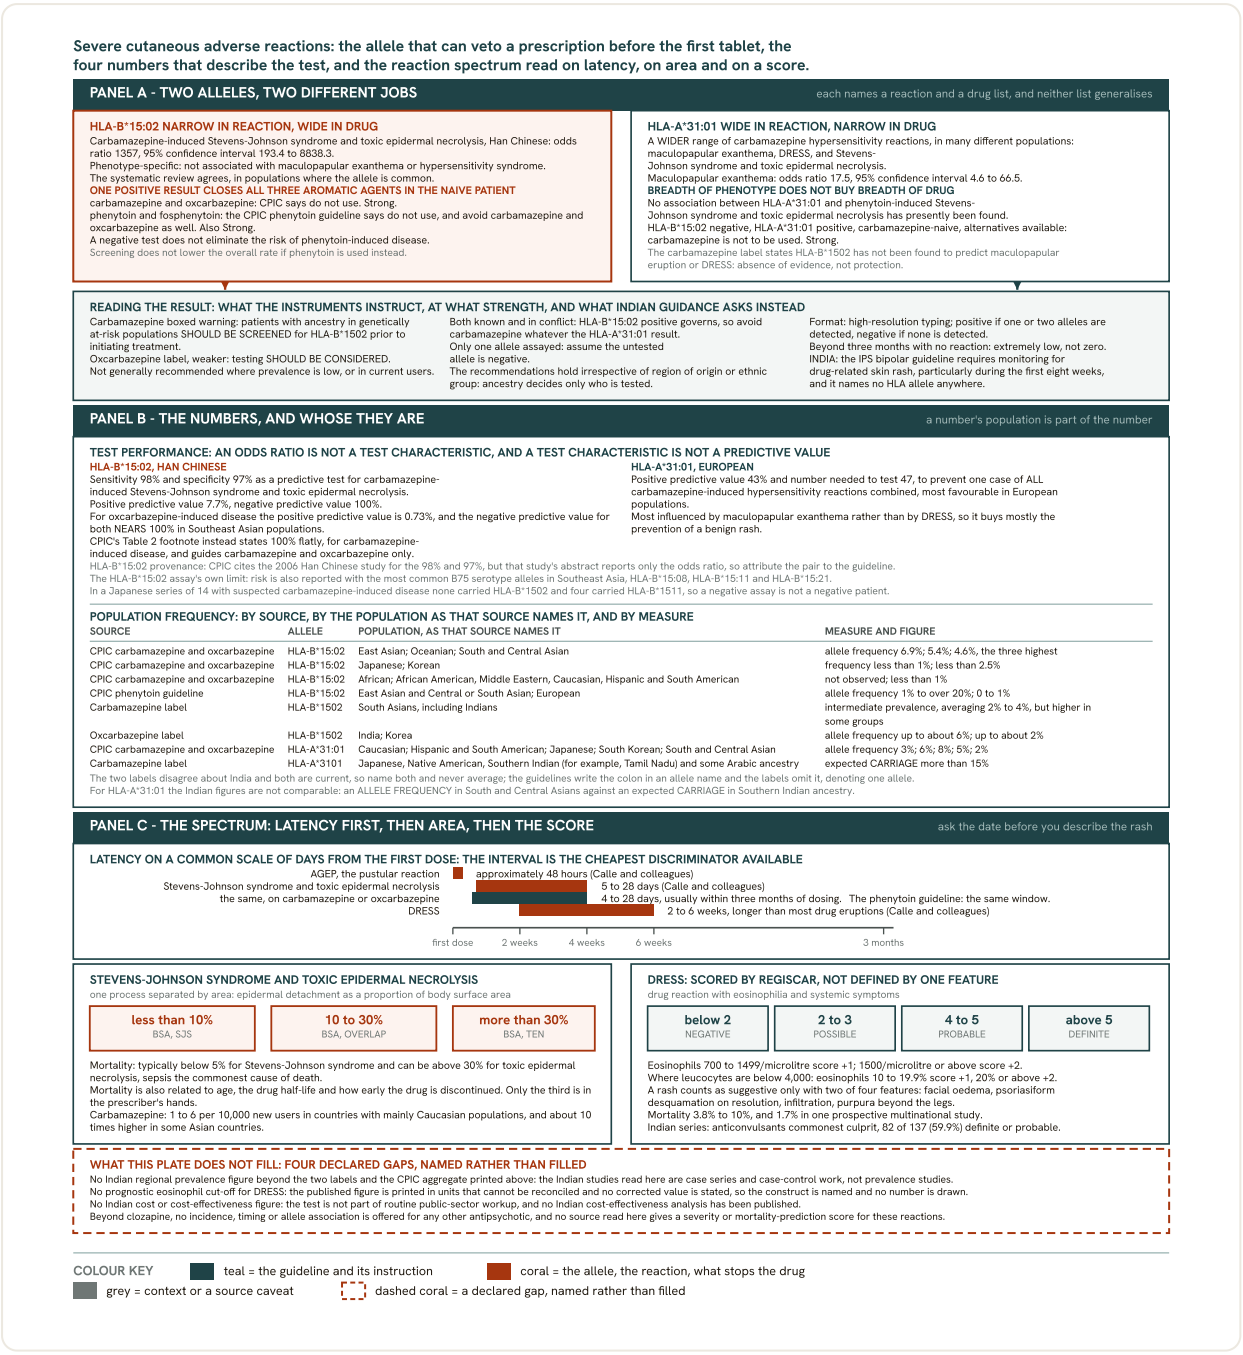


Fig 37.10 Severe cutaneous adverse reactions: allele to drug to population to test decision, then the spectrum the test exists to prevent. Two alleles do two different jobs. HLA-B*15:02 is narrow in reaction and wide in drug, carrying an odds ratio of 1357 for carbamazepine-induced Stevens-Johnson syndrome and toxic epidermal necrolysis in Han Chinese while being not associated with maculopapular exanthema or hypersensitivity syndrome, and one positive result closes carbamazepine, oxcarbazepine and phenytoin in the naive patient at Strong classification. HLA-A*31:01 is the mirror image, wide in reaction and narrow in drug, with no association with phenytoin-induced disease presently found. The four numbers answer four different questions and each carries its population: sensitivity 98% and specificity 97%, positive predictive value 7.7% against 0.73% for oxcarbazepine-induced disease, and a negative predictive value CPIC states twice and differently, 100% in its Table 2 footnote and nears 100% in its narrative, printed here as both rather than averaged. The frequency table is set out by source and by measure because the two labels disagree about India, 2% to 4% against up to about 6%, and because an allele frequency of 2% in South and Central Asians and an expected carriage above 15% in Southern Indian ancestry are two measurements and never one. The spectrum is then read on latency first, on body surface area second and on a score third, and Indian guidance asks for rash monitoring through the first eight weeks rather than a genotype. (CPIC carbamazepine and oxcarbazepine 2018 update; CPIC phenytoin 2020 update; the FDA carbamazepine and oxcarbazepine labels; Bastuji-Garin and colleagues 1993 for the body-

surface-area bands; Hung and colleagues 2006; Amstutz and colleagues 2014; Calle and colleagues 2023; Sasidharanpillai and colleagues 2022; IPS Clinical Practice Guideline, bipolar disorder, 2025.)

Strength of wording is the substance here, and flattening the instruments into one sentence about recommended testing throws away most of the marks. The carbamazepine label carries the strongest formal mandate: its boxed warning states that patients with ancestry in genetically at-risk populations should be screened for HLA-B*1502 prior to initiating treatment. The oxcarbazepine label is deliberately weaker, stating that testing for the presence of the HLA-B*1502 allele should be considered in patients with ancestry in genetically at-risk populations, prior to initiating treatment. It fences the recommendation from both ends: screening is not generally recommended in patients from populations in which prevalence is low, or in current users.

The test format is specified rather than left open. The carbamazepine label calls for **high-resolution HLA-B*1502 typing** and defines the result binarily, positive if either one or two alleles are detected and negative if none is. There is no intermediate genotype, so heterozygosity is not a partial result and does not license a partial decision.

The guideline recommendations, read against those labels, are more decisive and are graded.

- In a carbamazepine-naive or oxcarbazepine-naive patient who is positive for the allele, CPIC states that carbamazepine and oxcarbazepine should be avoided respectively, because of the greater risk of Stevens-Johnson syndrome and toxic epidermal necrolysis, classified Strong for both drugs.
- The phenytoin guideline extends the same logic across the class: in a phenytoin-naive patient positive for the allele, do not use phenytoin or fosphenytoin, and avoid carbamazepine and oxcarbazepine, also classified Strong.
- Where the second allele is positive and the first negative, CPIC's Table 2 states that for a carbamazepine-naive patient with alternative agents available, carbamazepine is not to be used, classified Strong.
- Where both results are known and conflict, the first allele governs: if the patient is carbamazepine-naive and positive for HLA-B*15:02, carbamazepine should be avoided regardless of the HLA-A*31:01 result.
- Where only one allele has been assayed, CPIC directs the clinician to assume the untested allele is negative rather than to treat the missing result as an unknown risk.

One further sentence disposes of the commonest abuse of this reasoning. CPIC records that these recommendations hold irrespective of the patient's region of origin or ethnic group. Ancestry decides who is offered the test; once a result exists it decides nothing further, and a positive result in a low-frequency population is still positive.

24.6 Timing: when the risk sits, and the patient already taking the drug

The risk is front-loaded, and every timing statement in this topic follows from that. CPIC states that the latency period for carbamazepine-induced or oxcarbazepine-induced Stevens-Johnson syndrome and toxic epidermal necrolysis is short with continuous dosing and adherence to

therapy, about 4 to 28 days, with cases usually occurring within three months of dosing; the phenytoin guideline prints the same window of 4 to 28 days for phenytoin-induced disease, again with cases usually occurring within three months of dosing. The carbamazepine label puts it as a proportion: over 90% of treated patients who will experience the reaction have it within the first few months of treatment.

That front-loading makes the treatment-experienced patient a different problem. CPIC states that patients who have been continuously taking carbamazepine or oxcarbazepine for longer than three months without developing cutaneous reactions are at extremely low risk, but not zero. Read the parenthesis: the tolerant patient is reassured, not discharged from vigilance, and an answer that writes off residual risk as nil has overstated a guideline that took care not to.

PITFALL

Treating an early rash two or three weeks into carbamazepine as a nuisance to be watched. It sits inside the stated latency window, and because mortality is related to how early the drug is discontinued, the decision to observe rather than stop is itself a determinant of outcome. The opposite error is stopping a drug taken uneventfully for years on the strength of a newly reported allele.

24.7 What a negative test does not buy

A screening programme is only as good as the drug that replaces the one it prohibits. CPIC records plainly that a HLA-B*15:02 screening policy for carbamazepine will not decrease the overall rate of Stevens-Johnson syndrome and toxic epidermal necrolysis if other anticonvulsants associated with the reaction, such as phenytoin, are used instead. The regulatory position agrees: consideration should be given to avoiding phenytoin as an alternative for carbamazepine in patients positive for HLA-B*15:02, because of the increased risk in patients of Asian ancestry.

TRAP **Substituting phenytoin for carbamazepine in a carrier.** It is the substitution a candidate reaches for first, and it is the wrong answer to write. One positive result rules out all three aromatic agents in the naive patient, and a screening policy that diverts patients from one to another prevents nothing.

Nor is the wider class clean. CPIC names **aromatic anticonvulsants** including eslicarbazepine, lamotrigine, phenytoin, fosphenytoin and phenobarbital as having very limited evidence, if any, linking them to Stevens-Johnson syndrome and toxic epidermal necrolysis through this allele, and states that caution should still be used when choosing an alternative agent. Cross-reactivity is also clinical rather than purely genetic: the carbamazepine label tells prescribers to inform patients that **about a third of those who have had hypersensitivity reactions to carbamazepine also experience hypersensitivity reactions with oxcarbazepine**.

The last limitation is in the assay itself. CPIC records that, in addition to the tested allele, risk for carbamazepine-induced Stevens-Johnson syndrome and toxic epidermal necrolysis has been reported in association with the most common B75 serotype alleles in Southeast Asia, HLA-B*15:08, HLA-B*15:11 and HLA-B*15:21. A Japanese series makes the consequence concrete:

among 14 patients with suspected carbamazepine-induced disease, none carried HLA-B*1502 and four carried HLA-B*1511. That is a case-series proportion without controls rather than a risk estimate, but it shows how a single-allele assay can return negative in a susceptible patient.

24.8 The Indian evidence base, and what Indian guidance actually says

One prospective study turned screening into policy anywhere. In the Taiwanese screening cohort, 7.7% tested positive and were advised not to take carbamazepine, being given an alternative or continuing their prestudy medication, while the 92.3% who tested negative were advised to take it; no Stevens-Johnson syndrome or toxic epidermal necrolysis developed in the screened-negative subjects receiving carbamazepine, against approximately 10 cases expected from the estimated historical incidence. The caveats should be stated: the comparison is against historical controls, and the expected cases are counted among study subjects rather than among the negatives who took the drug. There is no Indian equivalent of that trial.

What India has instead is three small studies, and their limits are the point. The foundational series found the allele in six of eight Indian patients with carbamazepine-induced Stevens-Johnson syndrome. A North Indian case-control study found the allele in 2 of 9 (22.2%) carbamazepine cases and in none of 37 carbamazepine-tolerant controls, p equal to 0.035, and in none of the 8 phenytoin cases. A South Indian study concluded that the association of the allele with phenytoin-induced Stevens-Johnson syndrome and toxic epidermal necrolysis is rare in the population studied, with small sample size acknowledged as its limitation. None is a prevalence study, and none supports a statement about how common the allele is in any Indian region. A systematic review places India inside the affected geography: allele-positive patients with carbamazepine-induced disease have been reported from Asian countries only, including China, Thailand, Malaysia and India.

IN INDIA

Indian ancestry sits inside the at-risk group the carbamazepine boxed warning names, so a genotype before the first tablet is the decision this evidence supports in a carbamazepine-naive or oxcarbazepine-naive patient. Where only one allele has been assayed, the guideline instruction is to treat the untested allele as negative and act on the result you have. A positive result closes off phenytoin as the substitute as well, which is the part most likely to be missed wherever phenytoin is the default second choice. Where no genotype can be obtained, the eight-week rash watch below is what remains, and it is not a formality.

Indian guidance has not followed the pharmacogenetic route. The Indian Psychiatric Society bipolar guideline states that patients taking carbamazepine must be monitored for drug-related skin rash, particularly during the first eight weeks of therapy, as the treatment is associated with life-threatening rashes, and it names no HLA allele anywhere. Its answer to this risk is clinical vigilance over a defined window rather than genotyping. The corresponding position in the schizophrenia guideline is equally cautious in general terms, holding that the clinical utility of psychopharmacogenomics is disputable because of the lack of larger standardised trials and cost-effectiveness evaluations. A candidate asked what Indian guidance requires should answer the

rash watch, and should not attribute a genotyping mandate to a document that does not contain one.

MNEMONIC

Allele, Agent, Ancestry, Already taken

The four questions that settle a stem: which allele is tested, which aromatic agent is being started, whether ancestry places the patient in the at-risk group the label names, and whether they are naive or have taken the drug uneventfully beyond three months. The fourth changes the answer more often than candidates expect.

A negative result narrows the risk and never abolishes it; a positive result forbids all three aromatic agents in the naive patient, whatever their ancestry.

25 · DRESS syndrome and other severe idiosyncratic reactions to anticonvulsants and antipsychotics

Most drug rashes declare themselves within a fortnight. DRESS often does not, and that single property is what makes it dangerous at the bedside and reliable in an examination. **DRESS**, drug reaction with eosinophilia and systemic symptoms, begins weeks after a drug was started, damages organs as readily as skin, and in Indian series is caused mostly by agents that sit on a psychiatric prescription. The marks are in three places: the time course that shifts the differential towards DRESS and away from the other severe cutaneous reactions, the scoring system an examiner will name and expect reproduced as reasoning rather than as a form, and the culprit drugs. Candidates lose them by quoting a diagnostic band as though it were a prognostic one, by carrying a figure out of the population it was measured in, and by waiting for a rash before believing the diagnosis.

The allele-testing framework and the severe cutaneous reaction spectrum in its own right are taught separately in these notes, and carbamazepine's own pharmacology belongs to the Psychopharmacology I: principles and core drug classes notes. What follows is the reaction itself: how it announces itself, how it is scored, which drugs produce it, and what can and cannot be said once the drug has been stopped.

25.1 The latency that names the syndrome

Calle and colleagues, in their World Allergy Organization Journal review, state that in most patients the reaction occurs **2 to 6 weeks** after the drug is started, a latency period longer than that of most drug eruptions. That property does a great deal of the diagnostic work before any morphology is described. A febrile eruption on the third afternoon of a new drug is usually something else; a febrile eruption in the second month, when the prescription has stopped feeling new and the drug chart has stopped being re-read, is the one that gets attributed to sepsis, to a viral illness, or to the psychiatric disorder itself.

The same review supplies the contrast that makes this examinable. Stevens-Johnson syndrome and toxic epidermal necrolysis run a shorter course from exposure, 5 to 28 days, and the pustular reaction **AGEP** shorter still, at approximately 48 hours. Latency shifts the differential but does not separate these three reactions by itself: the Stevens-Johnson window runs out to 28 days while the DRESS window opens at two weeks, so the two overlap for a fortnight, and morphology, mucosal involvement, systemic findings and investigations remain necessary. It is still the reason the first question in a suspected drug reaction is a date and not a description.

RULE **Establish when the drug was started before you characterise the rash.** Morphology overlaps across the severe cutaneous reactions and is read differently by different observers, whereas the interval between the first dose and the first symptom is a fixed number that the patient or the file can supply. In a syndrome whose whole difficulty is that nobody suspects a drug started six weeks ago, the interval is the cheapest discriminator available.

25.2 How common it is, and how badly it ends

Every published incidence figure for this syndrome carries a different denominator, and the three in common circulation are not interchangeable.

more than 1 per 10,000	0.9 per 100,000	2.18 to 40 per 100,000
MEDICATION EXPOSURES	INHABITANTS	INPATIENTS

Read these as estimates answering three different questions, not as one rate quoted three ways. An exposure denominator asks how risky a prescription is; a population denominator asks how much of this a health system will see; an inpatient denominator is enriched for polypharmacy, for severity and for the surveillance that produces a diagnosis at all, which is why its upper end sits far above the population figure. A syndrome that is under-recognised is also under-counted, so each estimate may be shifted by how reliably the diagnosis is made and recorded in the setting it was measured in, and none of the three is a clean measurement of risk.

Mortality is reported by the same group as 3.8% to 10% despite treatment, with 1.7% in one prospective multinational study. A single figure is therefore hard to defend: the prospective cohort and the pooled literature are not measuring the same population, and an answer that quotes only the low end understates a syndrome that kills, while one that quotes only the high end inflates it. Give the spread and attribute each end.

What cannot be offered here is a prognostic laboratory threshold. The eosinophil count said to mark a poor outcome is printed in the source literature in units that cannot be reconciled with the quantity they describe, and no corrected value is stated anywhere in that paper. Reconstructing the intended number would be an invention. Substituting the diagnostic eosinophil band described below would be worse, because a cut-off built to classify a case was never validated to predict its outcome.

RULE **A diagnostic threshold is not a prognostic threshold.** The two are derived from different questions on different samples. Where a prognostic cut-off is genuinely unavailable, the correct answer names the construct and says the number is not established, and prognosis is then discussed through mortality and through the pattern of organ involvement rather than through a laboratory line that does not exist.

25.3 The phenotype, and the warning that precedes the rash

The FDA-approved label for carbamazepine describes the reaction as presenting, typically though not exclusively, with fever, rash, lymphadenopathy and/or facial swelling in association with involvement of other organ systems, with eosinophilia often present. Four points in that sentence repay attention: fever is part of the syndrome and not an incidental infection; facial swelling is a positive feature rather than a variant of urticaria; the organ involvement is the reason the syndrome is severe; and the phrase "not exclusively" is the label refusing to make any single feature necessary.

The same label carries the warning that changes practice, which is that the early manifestations of hypersensitivity can be present before any rash is evident. A patient on a recently started anticonvulsant with unexplained fever and lymphadenopathy is already inside the differential. Waiting for skin before considering the drug converts a recognisable prodrome into a multiorgan presentation.

One negative statement from the same label needs teaching exactly as written. It states that the HLA-B*1502 allele has not been found to predict the risk of less severe cutaneous reactions such as maculopapular eruption, nor to predict DRESS. That is an absence of demonstrated predictive value, not a demonstration that carriers are protected: it means the allele cannot be used to anticipate this syndrome, and equally that a negative result cannot be used to dismiss it. Allele testing in its own right, with its indications and its operating characteristics, is taught elsewhere in these notes.

25.4 Scoring it: RegiSCAR as an examination construct

The **Bocquet criteria** came first and were simple enough to memorise: a drug rash, a haematological abnormality in the form of eosinophilia or atypical lymphocytes, either of which satisfies that element, and systemic manifestations. They define the syndrome but stop short of grading confidence in it. **RegiSCAR** was built to fill that gap. RegiSCAR sorts a suspected case by total score into four bands: **below 2 negative, 2 to 3 possible, 4 to 5 probable, and above 5 definite**. The examination point is that the instrument returns a level of certainty rather than a yes or a no, so a case can be reported as probable and treated as real.

Three of its parameters are worth knowing in detail because each is defined more tightly than candidates assume. Eosinophilia is scored in two bands, the higher scoring double: eosinophils 700 to 1499/microlitre earns one point and **1500/microlitre or above** earns two. Because a septic or marrow-suppressed patient may have a low total count, there is a percentage route as well: where leucocytes are below 4,000, eosinophils of 10 to 19.9% earn one point and 20% or above earn two. Lymphadenopathy is scored only when it clears a size and a site threshold,

namely nodes of more than 1 cm affecting two anatomical regions, which is why a single palpable cervical node adds nothing.

The rash parameter is the one most often misquoted. A rash counts as suggestive of the syndrome only when it carries at least two of four specified features: facial oedema, resolution with psoriasiform desquamation, infiltrated skin lesions, and purpuric lesions involving areas other than the legs.

MNEMONIC

FIPP: Facial oedema, Infiltrated lesions, Purpura beyond the legs, Psoriasiform desquamation on resolution

Two of these four are what turns a nonspecific exanthem into a rash the scoring system will accept as suggestive. The last is retrospective, which is a useful reminder that some of the score can only be completed as the patient recovers.

A further point is available for the exclusion of alternative causes. The footnote to the Indian scoring table adds one point if at least three of four listed tests were performed and found negative, the four being:

- antinuclear antibody
- hepatitis A, B and C
- infection due to Mycoplasma or Chlamydia
- blood culture

TRAP **Quoting the exclusion-of-alternatives item as though one agreed list existed.** The Indian paper's footnote enumerates the four investigations above. The World Allergy Organization Journal review states the same item differently and does not enumerate the same set. Neither wording is an error and they must not be blended into a single averaged list. Name the source you are quoting, and if the stem does not specify one, say that published versions of the item differ.

25.5 RegiSCAR against the Japanese criteria

The competing construct is **DiHS**, drug-induced hypersensitivity syndrome, for which the Japanese consensus group laid down seven mandatory features, one of them reactivation of HHV-6. Because that assay is unavailable in many centres, the same group proposed a modification with five essential features for atypical DiHS. The two systems are not two names for one thing, and the difference is examined through the parameters they treat as mandatory.

PARAMETER	REGISCAR	JAPANESE DIHS CRITERIA
Fever	No positive points; its absence instead scores one negative point	Mandatory for typical and atypical disease
Maculopapular rash	No positive points as one of the three features mandatory for DiHS	Mandatory for typical and atypical disease
Persistence of symptoms after the drug is stopped	No positive points	Mandatory for typical and atypical disease
Early resolution	No positive points, and resolution before 15 days scores one negative point	Not specified in the sources drawn on here
Viral reactivation	Not among the parameters described here	Reactivation of HHV-6 is one of the seven mandatory features; a five-feature version exists where testing is unavailable
Yield in an Indian tertiary centre, severe reactions	57 of 58 (98.3%) classified as definite or probable	32 of 58 (55.2%) classified as atypical DiHS

Two readings of that comparison are needed. First, Sasidharanpillai and colleagues make the statement about the three DiHS-mandatory parameters narrowly, and it should be read narrowly: it says those three earn no positive points as the mandatory triad, not that RegiSCAR never scores a rash at all. Second, the difference in yield in that Indian series was statistically significant at P less than 0.001, so the choice of instrument changes how many severe reactions are captured as cases, which matters for pharmacovigilance reporting as much as for the individual patient.

TRAP **Treating fever as a scoring point in RegiSCAR.** Fever is mandatory in the Japanese criteria and earns nothing positive in RegiSCAR, where its absence subtracts instead. A candidate who has memorised one system and answers a question set on the other will score the same patient two different ways and defend the wrong one. Say which instrument you are using in the first line of the answer.

25.6 The culprit drugs a psychiatrist actually prescribes

In the Indian tertiary-centre series, anticonvulsants were the commonest culprit class by a wide margin, accounting for 82 of 137 (59.9%) cases of definite or probable disease and 36 of 53 (67.9%) of the definite cases. That profile lands squarely on the mood-stabiliser prescription.

IN INDIA

In a patient presenting with fever, an eruption and deranged organ function, work the anticonvulsant column of the drug chart first rather than last, and interrogate everything started in the preceding weeks rather than only the most recent addition. In the Indian series above, the anticonvulsants supplied the majority of definite and probable cases, and the latency reported by Calle and colleagues means the offending drug will often have been started several weeks earlier, commonly within the preceding 2 to 6 weeks, at a point where both patient and clinician have stopped regarding it as new.

Lamotrigine is the agent whose risk a psychiatrist is most often asked to quantify. Its label reports serious rashes, including Stevens-Johnson syndrome, at approximately 0.3% to 0.8% in paediatric patients aged 2 to 17 years and 0.08% to 0.3% in adults. The timing follows the pattern of this whole section: nearly all life-threatening rashes have occurred within 2 to 8 weeks of treatment initiation, with isolated cases after prolonged treatment, for example at 6 months. Note that this window is drawn from a different source and a different endpoint than the general latency quoted earlier, so the two figures should be attributed rather than merged.

The interaction that matters most in bipolar practice is with valproate. In paediatric epilepsy trials, serious rash affected 1.2% (6 of 482) of children taking concomitant valproate against 0.6% (6 of 952) not taking it. The label also lists the presence of the HLA-B*1502 allele among the rash risk factors in its boxed warning, and reports that retrospective case-control studies in patients of certain Asian ancestry suggest an approximately 2 to 3 times higher risk of Stevens-Johnson syndrome and toxic epidermal necrolysis in carriers using lamotrigine. Lamotrigine is therefore not a self-evidently safe substitute in a carrier, and an answer that offers it as one should say which source it is following.

PITFALL

Quoting the valproate figures above when counselling an adult with bipolar disorder about a lamotrigine and valproate combination. Both come from paediatric epilepsy trials in children on multidrug regimens. No bipolar-population figure for that combination is offered here, and the honest counselling position is that the risk is real and quantified only in a different population.

25.7 Beyond DRESS: the other severe idiosyncratic reactions

Carbamazepine carries a second idiosyncratic reaction of boxed-warning severity that is haematological rather than cutaneous. Its label puts the risk of aplastic anaemia and **agranulocytosis** at **5 to 8 times** that of the general population, while noting that the untreated general-population risk is low, at approximately 6 patients per million population per year for agranulocytosis and 2 patients per million population per year for aplastic anaemia. Both halves of that warning must be taught together. A multiplier of that size applied to a base of that size still produces a rare event, which is why the reaction is not a reason to withhold the drug, and why the monitoring that accompanies it is justified by severity rather than by frequency.

The same label also describes a cutaneous pattern that is neither DRESS nor Stevens-Johnson syndrome. A fixed drug eruption presents as single or multiple circumscribed patches recurring at the same site with each exposure, and with repeated exposure the lesions may progress to a **generalised bullous fixed drug eruption**, a serious reaction that may resemble Stevens-Johnson syndrome or toxic epidermal necrolysis. The teaching value is the recurrence at a fixed site, a feature no other reaction in this group shares and one that is available from the history alone.

25.8 Clozapine, and the limits of the antipsychotic evidence

The antipsychotic side of this topic is smaller than it looks. A screen of the European pharmacovigilance database identified 51 probable or definite cases among 75,190 clozapine-treated patients with reports, with a median of 25 days between clozapine initiation and symptom onset and 3.9% ending fatally. A separate systematic review of the published literature found 26 reports covering 27 patients meeting RegiSCAR criteria, which is a collection of published cases and not an incidence cohort.

Those denominators must not be converted into a rate. A pharmacovigilance database counts reports, not patients at risk, and its denominator is the population for whom some report exists rather than the population exposed. Dividing the case count by that denominator produces a number that looks like an incidence and is not one, and because both the cases counted and the population reported on are shaped by who files a report, there is no defensible direction in which to correct it. Read the two studies for what they do establish: the reaction occurs on clozapine, the median interval from initiation to symptom onset in that pharmacovigilance series was 25 days, and a small proportion of reported cases die.

Beyond clozapine, the evidence assembled here stops. No incidence, timing or genetic-association data are offered for the other antipsychotics, and no allele association is stated for clozapine-induced blood dyscrasias or for clozapine-related DRESS. That is a limit of the sources and not a statement that the risk is absent, so the class-level generalisation is precisely what a careful answer refuses to make. The clozapine plasma level, its target range and the haematological monitoring schedule are taught separately under therapeutic drug monitoring, and the resistance decision that put the patient on clozapine is a separate question again.

25.9 After the reaction: confirming the culprit

Where several drugs were started together, the suspected culprit is identified at the bedside during the illness, and testing after recovery refines that attribution rather than making it for the first time. What follows the suspicion acutely is one step, and only one, that this chapter can source: the same authors who set the patch-testing interval, reviewing the French treatment algorithm, place interruption of the suspected drug inside the treatment of even the least severe scenario, alongside topical corticosteroids, emollients and antihistamines, and they describe the rash and the organ damage as resolving gradually after the involved drug is stopped. Withdrawal is therefore an acute treatment step and not something deferred until attribution is confirmed. Everything else about treatment stays untaught here: no corticosteroid indication, dose, route or taper, no second-line immunosuppressant, and no rule for what to do when several drugs remain

equally plausible. Calle and colleagues also state that **patch testing** should be performed 4 to 6 weeks after the reaction, with the patient off immunosuppressive therapy or systemic corticosteroids for at least 4 weeks, in order to reduce false negative results. The washout requirement is itself informative: it is a reminder that many of these patients will have been treated with systemic corticosteroids, and that testing them too early answers a question about their immunosuppression rather than about their drug.

GOOD TO REMEMBER

Do not send the patient for culprit testing at the discharge visit. Testing performed inside the washout period reads negative for the wrong reason, and a false negative here is more harmful than no test at all, because it returns the patient to a drug that nearly killed them with a piece of paper saying it did not. Book the test after the interval and after the steroids have stopped, and record the suspected drug in the chart in the meantime.

One Indian contribution is worth knowing as depth rather than as standard practice. A prospective Indian validation study of the **CET score** as a bedside screen reported that, at a cut-off above 2.12, it achieved sensitivity 94.3%, specificity 60%, positive predictive value 80.5% and negative predictive value 85.7%, against a median diagnostic delay using RegiSCAR of around 14.5 hours. Read those operating characteristics correctly: high sensitivity with modest specificity is the signature of a screening instrument, so it is a tool for deciding whom to work up quickly, not a replacement for the scoring system that classifies the case.

Fever with eosinophilia and deranged organ function weeks after a new anticonvulsant is DRESS until another cause is proven, and the rash may not have arrived yet.

Consolidate and apply

26 · Worked vignettes

A vignette in this territory is never really about the drug; it is about what changed around the drug. Stems that carry marks here rarely open with an undiagnosed patient. They open with someone already on treatment, then introduce one further fact: a second prescription, a failing organ, a pregnancy, an inherited metabolic capacity, a concentration somebody has measured, a trial that did not work, or a syndrome the treatment itself produced. That further fact is the **second variable**, and the examination value of the stem lies entirely in whether the candidate reasons from it or reads past it. This closing section introduces no new agent and no new class; it rehearses the reasoning assembled across this chapter as a reading skill, working each stem from the second variable inwards.

The discipline that earns marks is refusing to answer with a treatment algorithm. A patient who destabilised after a physician added something is not asking which antidepressant treats depression; it is asking which clearance pathway two drugs share. A patient whose measured concentration comes back low is not asking whether the drug works; it is asking whether it was

taken, whether the sample means anything, and what a low value does to the claim that the illness is resistant. Where a stem genuinely turns on class pharmacology itself, the receptor logic behind an adverse-effect signature or a choice within a single class, that reasoning belongs to the Psychopharmacology I: principles and core drug classes notes. What this chapter supplies is the layer sitting above those decisions.

26.1 Four questions locate the second variable

Before deciding what a stem is asking, run the same four questions across it. Between them they cover almost every way an examiner can complicate an ordinary prescription inside one paragraph, and each opens onto a different section of this chapter.

Company	Organ	Genome	Sequence
WHAT ELSE IS ON BOARD	WHAT IS CLEARING IT	WHAT THE PATIENT BROUGHT	WHAT HAS ALREADY HAPPENED

Company asks what else the patient is taking, including everything no psychiatrist prescribed: the physician's additions, the over-the-counter purchase, the traditional remedy, tobacco. Organ asks which tissue transforms and excretes this agent, and whether it is currently doing that work. Genome asks what metabolic capacity and what risk alleles the patient brought into the room. Sequence asks what has already been given, at what exposure, for how long, with what adherence, and what it produced. The fact a stem plants to make one of the four decisive is its discriminator, and a well-built stem carries exactly one.

RULE **The prescription is the setting; the second variable is the question.** If your answer can be written without naming the interaction, the organ, the allele, the measured concentration or the failed sequence, you have answered a question about the diagnosis rather than the one the stem asked. The diagnosis here is usually scenery: it tells you which class the patient is on and nothing more. The marks sit in naming the second variable, stating the mechanism by which it acts on this prescription, and choosing the move that addresses that mechanism.

26.2 The vignette that resolves on another drug

The commonest stem describes a patient stable for a long time who is not stable now, while the psychiatric prescription is unchanged. Somebody else changed something. The reading move is to date the deterioration against the whole medication list rather than the psychiatric part of it, then ask by what route the new agent could reach the old one. The routes behave differently. A **pharmacokinetic interaction** changes how much drug is present, by altering absorption, protein binding, metabolism or renal handling, and the picture is that of too much or too little of an agent whose dose nobody touched. A **pharmacodynamic interaction** changes what a drug does rather than how much of it is there, most often by two agents acting additively on one physiological endpoint; the concentration need not move for the interaction to be real, and one pair of drugs can sit on both axes at once.

Candidates instinctively look for the enzyme, and enzymes do account for many of these stems, but the interaction matrix this chapter develops is deliberately not closed at the cytochrome axis. Drugs also compete for renal excretion, alter the sodium handling that decides how much of a

renally cleared agent is retained, displace one another from plasma protein, and add their effects at a shared receptor or a shared electrophysiological target with no enzyme involved. A stem placing a diuretic, an anti-inflammatory or an antihypertensive beside a psychotropic is usually testing that non-enzymatic axis, and a candidate hunting for a cytochrome substrate will not find one. The agent-by-agent matrix is set out in this chapter's interaction section; what the vignette wants first is the route, named correctly.

GOOD TO REMEMBER

Ask the question a psychiatric review rarely asks: what did somebody else start or stop, and when. Take the physician's list, the surgical discharge summary, the over-the-counter purchases and the herbal preparations, and lay the onset of the new problem against them. Ask about tobacco too, because moving a patient onto a smoke-free ward alters an inducing exposure without a single prescription being rewritten. Anything lining up in time with a change outside the psychiatric list is an interaction until it has been reasoned away.

26.3 The vignette that resolves on an organ

A second family holds the drug list constant and changes the patient. The physiology has moved: renal function has fallen with age or with intercurrent illness, the liver is cirrhotic, cardiac output is poor, or a neurological disorder has altered the tissue the drug acts on. The reading move is to ask which organ this agent depends on, then whether the stem has quietly told you that organ is compromised. **Renal elimination** and **hepatic biotransformation** are not interchangeable answers: an agent excreted essentially unchanged behaves quite differently in kidney disease from one extensively metabolised first, and the two failures call for different adjustments.

The recurring organ stems are cardiac, hepatic and neurological. The cardiac stem asks not only whether an agent is tolerated in heart disease but what it does to conduction and to the metabolic profile that helped produce the disease, and this chapter's monitoring section carries both halves. The hepatic stem asks about reduced clearance, and about sedation a compromised liver tolerates poorly. The neurological stem is subtlest: whether the psychotropic makes convulsions more likely, whether it worsens the movement disorder already being treated, or whether the antiparkinsonian medicine and the antipsychotic are pulling in opposite directions at one transmitter system. The vignette's job is to make the candidate notice that the organ, not the diagnosis, decides the answer.

26.4 The vignette that resolves on a measured concentration

A number in a stem is an invitation to interrogate it, not to recite a range against it. When a stem supplies a plasma concentration, the testable content is what that value means, and its meaning depends on when the sample was taken relative to the dose and to steady state, on which assay produced it, and on the patient's adherence beforehand. A value drawn before the concentration has plateaued is a point on a rising curve, not the patient's exposure. A value drawn at the wrong point in the dosing interval is not comparable with a range derived from a standard sampling time. The ranges, the sampling times and the monitoring schedules are taught in this

chapter's therapeutic drug monitoring section, and a vignette answer reasons against them rather than reproducing them.

The payoff arrives when the value is low. A low concentration in a patient who is not improving has competing explanations: the dose is genuinely too small, the drug is being cleared faster than expected because of an interaction or an inherited phenotype, or the patient is not taking it. Helping to investigate **covert non-adherence** against true pharmacological failure is one of the things a concentration does better than clinical judgement alone, which is why a level belongs in the assessment of apparent resistance before the sequence is escalated. It does not settle the question by itself. A value of zero establishes only that the drug has not been taken recently, and a value above zero is compatible with erratic taking, with full adherence, and with long-standing non-adherence disguised by a few recent doses. What turns a level into an answer is the timing of the last dose against the draw, the dose history, the interacting agents and, where doubt persists, a repeated sample. A high value in a patient who looks well matters too: the margin has narrowed, and the next intercurrent illness will be less forgiving of it.

PITFALL

Acting on a concentration without recording when the previous dose was swallowed. The result travels through the notes as a bare number, and by the time somebody decides, nobody can say whether the sample sat near the peak or near the trough, or whether the patient had been taking the drug at all beforehand. The dose then moves on a value that never described the patient's exposure. Write the time of the last dose and the time of the draw on the request form, and treat a result carrying neither as uninterpretable rather than reassuring.

26.5 The vignette that resolves on the genome

A third family plants an inherited fact. Sometimes it is explicit, a reported metaboliser status or an ancestry; more often implicit, in a patient with marked adverse effects at an ordinary dose, or no effect at a dose that suits most people, or a severe rash soon after an anticonvulsant was started. The reading move is to separate two quite different genetic questions that examinations like to set side by side. A **metaboliser phenotype** is a quantitative statement about clearance: it predicts the exposure a standard dose will produce, and the answer it drives is an adjustment of dose or a move to an agent independent of that pathway. A **risk allele** for a severe cutaneous or systemic reaction is not quantitative at all: it predicts a categorical hazard no dose adjustment addresses, so the decision it drives is whether to prescribe that drug at all.

The discrimination follows from that difference. Testing a metaboliser phenotype is informative mostly after an unexpected response, or before an agent with a narrow margin whose clearance that pathway dominates. Testing a risk allele is informative before the first exposure, in populations where the allele is common enough for the result to change practice, and its worth depends on the balance between how many affected patients it catches and how many unaffected patients it labels. This chapter's pharmacogenomics section teaches the enzymes and their application; the allele testing framework, the drugs it covers and the severe reaction spectrum sit in their own section beside it.

26.6 The vignette that resolves on a failed sequence

The resistance stem gives a history rather than a moment: a list of agents tried, with doses and durations, and a patient still unwell. Candidates rush to the next drug. The mark is in auditing the sequence first. Before an illness is called resistant, each previous trial has to be examined for whether the dose reached an adequate exposure, whether it ran long enough for a response to be judged, whether the patient actually took it, and whether the diagnosis and any maintaining physical or substance-related factor were right. A sequence failing those tests is **pseudoresistance**, and treating it as true resistance escalates a patient who has not yet had one adequate trial. This is where the measured concentration and the resistance assessment meet: a subtherapeutic level inside a supposedly failed trial reclassifies that trial instead of confirming it.

Once the sequence is genuinely adequate, the options are structurally distinct rather than a menu of equals. Switching replaces the agent; combination runs a second agent of the same class alongside the first; augmentation adds an agent from outside that class, which works through a different mechanism while the original treatment continues unchanged. The distinction turns on the class the added agent comes from, not on whether that agent has a therapeutic role of its own in the disorder being treated, because several of the standard augmenting agents do. The definitions of resistance, the staging models, how many failed trials each requires and the comparative evidence for each strategy are developed in this chapter's resistance and augmentation sections. For most of the modalities positioned against resistance, the decision moves to the device-based treatments once the drug sequence has been exhausted, and their indications and utility across disorders are taught in full in this chapter; the procedural technique for each is set out in the ECT and neurostimulation notes. Electroconvulsive therapy may enter earlier, because it is not positioned behind an exhausted drug list where a rapid or definitive response is required, and a candidate who makes it wait for one has misread the indication; the indication itself, with the body that states it, sits in this chapter's ECT section.

TRAP **Calling an illness resistant on the strength of a drug list.** Candidates read a stem containing several agent names and answer with the resistance pathway, because a long list looks like evidence. It is not evidence until every entry has been audited for adequacy of dose, adequacy of duration, adherence and diagnostic accuracy. Examiners write those failures into the stem deliberately, most often as a dose that never rose, a trial abandoned early, or a patient who quietly stopped.

26.7 The vignette that resolves on a syndrome the prescription produced

The next family reverses the direction of reasoning. The stem describes a syndrome and expects the candidate to read it backwards to the prescription that produced it. A febrile, rigid, confused patient on psychotropic medication is an urgent diagnostic and pharmacological problem at once: review the prescription immediately, asking which agents on this list act on which transmitter system, what was added or increased recently, and which of the drug-induced syndromes that combination makes possible, while the competing medical causes are evaluated alongside. The drug chart narrows the field faster than anything else available at the bedside, but it does not license deferring the search for sepsis, for central nervous system disease or for a

severe catatonic state, and this chapter's syndrome sections treat the exclusion of those competing causes as part of making the diagnosis rather than as an afterthought to it. The **iatrogenic phenocopy** is the trap in both directions, because a drug-induced syndrome can be mistaken for deterioration of the illness and deterioration of the illness for toxicity, and the discriminating information comes from the drug history, the tempo of onset and a directed neuromuscular examination taken together, with no single axis decisive on its own. What cannot separate them is whatever the two pictures hold in common, however unwell the patient looks.

The decisions that follow are pharmacological. The first is what to stop. Where the syndrome is severe, where several agents are acting additively, or where the diagnosis is still unresolved between two drug-induced syndromes, every plausible causative agent is withheld rather than one culprit chosen and the rest left running; the list is narrowed afterwards, and only where narrowing it is clinically safe. Naming the agent that most plausibly drove the syndrome still earns the mark, but it is an account of the mechanism, not a licence to continue the others. The second is the **rechallenge decision**: whether the same agent can be reintroduced, whether another member of the class is a safer destination, and what is monitored afterwards. Those judgements, with the precipitant combinations, the mechanism, the named criteria sets as examination constructs and the pharmacological management, are taught in this chapter's sections on the individual syndromes and on their comparative differentiation. The emergency-department presentation, the resuscitation and critical-care escalation and the disposition decision belong to the Perinatal/women's mental health and emergency psychiatry notes, and an answer here stops cleanly at that boundary.

26.8 The vignette that resolves on a physiological state

Pregnancy, lactation, advanced age and childhood are second variables of the same species as an organ, except that they move several things at once and often carry a second patient or a developing one. The stem usually presents a woman established on treatment who is now pregnant, an older patient on several medicines who has become confused or unsteady, or a child in whom an adult regimen has been scaled by weight and something has gone wrong. Distribution and clearance are altered, target sensitivity changes, and the risk being balanced is no longer only the patient's own.

The reading move is to name what has changed and for whom, before naming any drug. In the older patient the variables are cumulative: reduced clearance, altered target sensitivity, more co-prescriptions and poorer tolerance of anticholinergic and sedative burden, which is why the answer there often lies in subtraction rather than substitution. In pregnancy and lactation the balance sits between exposure of the fetus or infant and the risk carried by an untreated maternal illness, so a decision to stop is itself a decision with risk on both sides. The agent-level detail across all four groups is developed in this chapter's lifespan section.

26.9 Look-alike stems and the answering spine

Consolidation is easiest in pairs. The table sets out the stems most often confused in this territory, against the second variable that separates them and the axis the reasoning lands on.

TWO STEMS THAT LOOK ALIKE	THE SECOND VARIABLE THAT SEPARATES THEM	THE AXIS IT LANDS ON
Deterioration in a patient whose psychiatric prescription never changed, against deterioration of the illness	Whether a non-psychiatric medicine was started or stopped, and when	Company
Toxicity on unchanged dosing, against toxicity after a dose increase	Whether renal or hepatic function has fallen since the dose was set	Organ
Marked adverse effects at an ordinary dose, against adverse effects from a co-prescribed inhibitor	Whether the reduced clearance was inherited or acquired	Genome against Company
A severe reaction to an anticonvulsant, against a benign drug eruption	Whether a categorical risk allele rather than an exposure level is what is being tested	Genome
A low measured concentration from too small a dose, against one from covert nonadherence	Whether exposure was ever achieved at any point in that trial	Sequence
A genuinely resistant illness, against a run of inadequate trials	Whether each earlier trial reached adequate dose, duration and adherence	Sequence
A drug-induced syndrome, against a deterioration of the disorder being treated	Whether the tempo, the drug history and the neuromuscular findings fit the prescription or the illness	Sequence

MNEMONIC

Company, Organ, Genome, Sequence

The four second variables spell COGS, and the prescription is only one cog: the vignette asks what the others are doing. **C**ompany, what else is on board, prescribed or not. **O**rgan, which tissue clears the agent and whether it is failing. **G**enome, what capacity and what risk alleles the patient brought. **S**equence, what has been tried, whether it was adequate, and what it produced. Each opens a different section of this chapter, and a stem is usually built on exactly one.

With the second variable identified, an answer that earns marks follows four moves in order.

- Name the second variable in pharmacological language, not as a description of the patient.
- State the mechanism by which it acts on this particular prescription.
- Choose the move that addresses that mechanism: a dose adjustment, a switch off the shared pathway, a measured concentration, a test before prescribing, or an audit of the sequence.
- Say what you would not do, and why the obvious algorithmic answer leaves the second variable untouched.

That last move is where the pharmacology becomes visible to an examiner. A candidate who answers that the drug should be stopped has answered a question about a patient. One who answers that it need not be stopped, because the problem is an added inhibitor that will clear, a level drawn at the wrong moment, or a trial that was never adequate, has answered a question about pharmacology. Every vignette in this territory rewards the second answer.

> Find the second variable - company, organ, genome or sequence - and answer with the mechanism by which it acts on this prescription, never with the next drug in the algorithm.

27 · Consolidation and quick review

Nothing in this chapter is about a drug on its own. Every section here starts where a psychotropic meets something else: another agent competing for the same clearing enzyme, an organ that no longer clears what it used to, an inherited difference in handling or in immune recognition, a life stage in which the exposure is shared, a treatment course in which adequate trials have already failed, or a syndrome the prescription itself produced. The pharmacology of the agent as an object was settled in the chapter before this one. What is settled here is what happens to that pharmacology once it stops being the only thing in the room.

Call that something else the **second variable**. This closing section teaches no new agent and no new syndrome, because each has been taught in the section that owns it. What remains is to state the argument those sections were arranged to make, and to convert it into something that can be produced from a blank page with the book shut. The argument is that safe psychopharmacology is the management of second variables, and that the reasoning is the same every time: work out what the second variable does to exposure or to effect, measure the thing that has changed where a measurement will inform the decision, and act on the established risk or the observed change rather than on the label the patient arrived with.

27.1 One argument, tested against six second variables

Read as a list of territories, this chapter is long and disconnected: emergencies, resistance, monitoring, interactions, prescribing in physical illness and across the lifespan, pharmacogenetics and the device-based treatments. Read as one argument it is a single claim tested six times. The claim is that a psychotropic's behaviour is predictable once you have named what else is acting on it, and that the things that can act on it are few and recurring.

The six recur because they are the routes this chapter's material keeps arriving by: something competing with the drug pharmacologically, a fall in the capacity to clear or tolerate it, a constitutional difference in how the molecule is handled or recognised, a life stage the standard dose was not derived in, a treatment history of adequate courses that already failed, and the drug having produced the presentation itself. Every item in this chapter is one of those six wearing a clinical name, which is why the same three moves answer all of them.

RULE **A psychotropic prescription is safe or unsafe only in the presence of its second variable, never on its own.** The agent's own pharmacology sets what it can do; the second variable sets what it will do in this patient. An answer that describes the drug and stops has described only what the molecule brings, and left out everything else acting on it.

27.2 Exposure or effect: the question that opens every answer

Before anything can be monitored or adjusted, the second variable has to be sorted into one of two kinds, because the two are managed differently and confused constantly. An **exposure change** means the same prescribed dose now produces a different concentration at the site of action: the drug is cleared faster or slower, absorbed or distributed differently, or displaced from the protein carrying it. An **effect change** means altered biological activity at a shared endpoint: two agents push that endpoint in the same direction, a receptor population has changed, or the tissue is more vulnerable than it was. Concentrations may be entirely unchanged in a purely pharmacodynamic interaction, but constancy of concentration is not part of the definition, and the two mechanisms coexist in the same pair often enough that neither excludes the other.

The distinction decides the response. An exposure problem is answered by changing exposure: the dose, the interval, the route, the co-prescribed inhibitor or inducer, or the agent itself if its elimination route is the compromised one. An effect problem is answered by reducing or removing one of the converging contributors, with direct monitoring of the shared endpoint, or by accepting the combination under closer observation; a level inside the reference range does not exclude benefit from dose reduction, because a concentration-dependent effect can still be lessened by lowering exposure. Candidates who cannot separate the two reduce a correct dose and leave an additive combination untouched.

It also decides whether a measurement will help at all. A concentration informs the exposure question but cannot by itself rule an effect problem in or out, so a level requested in a purely pharmacodynamic problem returns normal and is then misread as reassurance; the effect endpoint itself, clinical, electrocardiographic or laboratory, has to be watched directly.

27.3 The three moves, and why they run in that order

Once the second variable is named and sorted, the same sequence organises the answer whichever of the six it is. It is a revision scaffold rather than a prescribing algorithm, and it yields in two situations: act immediately where delay is unsafe, as in a suspected drug-produced emergency, and act on what is already established where a contraindication or a positive pretreatment test is known before the prescription. Measure first only where the result is needed to choose the action. Outside those, the order matters, because each step is meaningful only if the one before it has been done.

- Anticipate. Before the prescription is issued or changed, state what this second variable does to exposure or to effect. This is a prediction from pharmacology, and needs no measurement.
- Measure. Monitor the thing the prediction says will move, alongside the routine monitoring the drug and the patient already require, and expect more than one endpoint to need watching at once. Monitoring chosen this way is **anticipatory monitoring**: the test is selected

because a specific mechanism predicts a specific change, not because it belongs to a routine panel.

- Act. Change the prescription on the strength of either a sufficiently established prospective risk or an observed change, according to the urgency and the evidence, and not on the diagnostic label, the drug's reputation, or the fact that the patient has tolerated the agent before. A contraindicated interaction, organ impairment, a pregnancy hazard, a positive pretreatment test and a suspected toxic syndrome are all acted on before any injury has been observed.

Reversing any two produces a recognisable clinical error. Measuring without anticipating gives a number nobody knows how to read; acting without measuring leaves the prediction untested, which is defensible only where waiting would be unsafe or where the risk was already established; and anticipating and measuring without acting is the commonest of the three, which is how an abnormal result sits in a chart while the prescription that produced it continues unchanged. A fourth step closes the loop outward: reporting what happened, so that a hazard invisible to any single prescriber reaches the surveillance system that can see it.

Anticipate	Measure	Act	Report
WHAT THE VARIABLE CHANGES	THE THING PREDICTED TO MOVE	ON THE RISK OR THE CHANGE	SO THE NEXT PATIENT BENEFITS

27.4 A frame that names all six

Because the six recur and the three moves do not change, one retrieval frame serves the whole chapter. Hold it as six prompts rather than as six topics: under pressure the failure is almost never that a candidate does not know a second variable exists, but that they answer on the one the stem made obvious and never scan for the others.

MNEMONIC

Second drug, End-organ capacity, Constitution, Ontogeny, Non-response, Drug-produced syndrome

The six second variables spell the thing they are: **S**econd drug, meaning another agent competing for the same clearing enzyme, displacing the drug from protein binding, altering its renal handling, or pushing the same physiological endpoint in the same direction; **E**nd-organ capacity, meaning cardiac, hepatic, renal or neurological function below what the standard dose assumed; **C**onstitution, meaning the inherited metabolising phenotype and the immune risk alleles that predict a severe reaction to a specific agent; **O**ntogeny, meaning pregnancy, lactation, childhood and old age as physiological states rather than as diagnoses; **N**on-response, meaning adequate trials that have already failed; and **D**rug-produced syndrome, meaning the acute psychotropic emergencies and the severe idiosyncratic reactions. Scanning all six before answering is the habit this chapter exists to build.

Each of the six is taught in full in the section that owns it, with its own agents, mechanisms and figures. The frame below is a retrieval index rather than a teaching table, and it deliberately gives

the same four columns to all six, because the point of the chapter is that the columns do not change.

THE SECOND VARIABLE	WHAT IT CHIEFLY CHANGES	WHAT IS WATCHED	WHAT THE CHANGE MEANS FOR THE PRESCRIPTION
A second drug taken alongside	Exposure where the agents meet at an enzyme, transporter or binding protein; effect where they converge on one endpoint	The drug concentration for the exposure route where a validated and clinically useful assay exists, otherwise an appropriate clinical or laboratory exposure marker; the shared endpoint for the effect route	Separate the agents, change one, or adjust exposure. Not every interaction is enzymatic: some are renal, some purely additive
End-organ capacity that has fallen	Mostly exposure, through reduced clearance; for some agents the tolerable effect narrows as well	Organ function and organ injury markers, interpreted alongside a drug level only where a validated level exists	Prefer the agent whose elimination route is intact, and set the starting exposure to the organ rather than to the diagnosis
Constitution: the inherited handling and the risk alleles	Exposure for the metabolising variants; the probability of a specific severe reaction for the immune risk alleles	The phenotype where testing is available, ideally before the first prescription rather than after a reaction	An unexpected dose response is a signal to assess phenotype alongside adherence, interacting drugs, organ function and sampling validity, none of which excludes the others. Where the owning section indicates a pretreatment test, testing precedes prescribing
Ontogeny: pregnancy, lactation, the very young, the old	Exposure through altered physiology, and the number of people exposed to it	Both parties where the exposure is shared, and the endpoints the age group is vulnerable on	The dose derived in a healthy adult is a hypothesis, not a prescription, and the risk of leaving the illness untreated belongs in the same calculation
Non-response after adequate trials	Neither exposure nor effect, but the meaning of every subsequent prescription	Whether each trial was adequate in exposure and in duration, and whether the drug was actually taken	Establish adequacy before declaring resistance. A course never absorbed or never swallowed is not a failed trial of the drug
A syndrome the drug itself produced	The whole question, from what to add to what to withdraw	The temporal relation between the prescription change and the onset, read together with the syndrome-specific examination findings	Withdraw or withhold the responsible agent first. The rechallenge and substitution decision is a later and separate judgement

27.5 The monitoring layer, and what each measurement is really for

Monitoring is the chapter's connective tissue, and it is badly taught as a schedule to memorise. Monitoring detects drug exposure, the drug's own effects and the patient-specific modifiers acting on either, and much of what it detects is clinically silent, so each test is best held as a **monitoring proxy**: a visible quantity standing in for an invisible one. What a test proxies for tells you when it is worth taking, what a normal result fails to exclude, and which second variable an abnormal one points at.

WHAT IS MEASURED	WHAT IT IS ACTUALLY A PROXY FOR	THE SECOND VARIABLE IT EXPOSES FIRST
The drug's plasma concentration	Exposure achieved, as distinct from the dose written on the chart	A changed clearance, a competing agent, or a prescription that was never taken
The corrected QT interval on the electrocardiogram	A shared cardiac effect that several unrelated agents can each contribute to	Additive pharmacodynamic interaction, and patient-specific susceptibility to delayed repolarisation and torsades
Weight, glycaemia and the lipid profile	Cumulative metabolic effect, which accrues without any symptom announcing it	A slow effect burden that no acute presentation will reveal
The blood count	An emerging haematological adverse effect tied to particular agents, which the patient cannot feel and cannot report	A drug-produced reaction caught as it develops, rather than a constitutional susceptibility the count could have identified in advance
Renal function	The remaining capacity to eliminate what the kidney clears	An end-organ change that happened after the dose was set
The liver function tests	Hepatocellular injury. They detect damage and do not quantify metabolising capacity, so a normal set does not establish that clearance is intact	An end-organ change that happened after the dose was set, and sometimes a hazard the drug itself produced
The reported adverse reaction	A population signal that no individual clinic can generate alone	A hazard that is real in aggregate and invisible at any single bedside

Held this way, the monitoring sections stop being a timetable and become the second half of a prediction. It also supports **reverse retrieval**, the drill worth doing: instead of asking what to monitor for a given drug, take an abnormal result and produce the second variables that could have generated it. That is the direction an examiner's stem runs, and the direction a referral arrives in.

27.6 When the drug itself is the second variable

Two of the six deserve separating out, because in both the prescription has stopped being the intervention and become part of the problem. The first is non-response. Once adequate trials have failed, the question is no longer which agent is theoretically best but what the failure has

established, and the answer is to check exposure before accepting effect: a course that never reached an adequate concentration, or was never taken, teaches nothing about the drug. Only once adequacy is verified does the switch, combine and augment decision become real, and the same verification is what positions the resistance-indicated device treatments taught here as modalities in their own right. That positioning is not universal: electroconvulsive therapy may be used first line where a rapid or definitive response is required, and is not conditional on failed adequate drug trials.

The second is the drug-produced syndrome, where the discipline is the **iatrogenic differential**: asking of any new presentation in a patient on psychotropics whether the prescription could have caused it, before asking what else the patient has developed. Rigidity, fever, confusion, a rash, a fall in the blood count and a disturbance of consciousness are all things a psychiatric patient can acquire independently and all things a psychotropic can produce. Timing against the last prescription change, the exposure history and the syndrome-specific examination findings have to be integrated, and none of them is universally sufficient on its own, which is what makes the drug history part of the investigation rather than its background. What this chapter owns is the pharmacology: which agents and combinations precipitate each syndrome, the mechanism, the monitoring that anticipates it and the pharmacological management. The emergency-department presentation, the resuscitation and critical-care escalation and the disposition decision are set out in the Perinatal/women's mental health and emergency psychiatry notes.

27.7 The two ways candidates lose these marks

Two failures account for most of the marks lost here. The first is **label-driven prescribing**: answering from the diagnosis and the drug class when the stem's real content is the second variable it quietly supplied, whether that is a physical illness, a co-prescription, an ancestry, a pregnancy or a previous failed course. The second is treating monitoring as a ritual, reciting what is checked without saying what it is checked for, which reads as memorised and does not survive a supplementary question.

TRAP **Answering the drug question the stem appears to ask instead of the second-variable question it actually asks.** Vignettes here are built by taking an ordinary prescribing decision and adding one complication, so the marks sit almost entirely in that complication. Candidates who begin by classifying the drug and end by naming an agent write a fluent answer to the wrong question; an answer that opens by naming the second variable and saying whether it changes exposure or effect has already earned most of what is available.

PITFALL

Continuing an unchanged prescription through a changed patient. The dose was correct for the renal function, co-prescriptions, weight and physiological state that existed when it was written, and none of those is fixed. Dehydration, a medicine started elsewhere, a pregnancy, an intercurrent illness and simple ageing all move the second variable while the chart stays still. Reviewing the prescription against the patient's current state, rather than reissuing what was tolerated before, is the most transferable habit here.

27.8 Where this chapter stops, and what to carry in

A disciplined answer is defined as much by what it hands over as by what it teaches.

Consolidating this material includes consolidating its edges, because several neighbouring chapters treat the same drugs and the same syndromes from a different angle. Each of the following is worth rehearsing as a **boundary handover**: named where it touches the question, passed to its owner in one sentence, and left there.

- The principles this chapter applies, meaning half-life and steady state, the cytochrome substrate map, the therapeutic index and choice between agents within a single class, belong to the Psychopharmacology I: principles and core drug classes notes. This chapter uses them as instruments rather than deriving them.
- The emergency act for the drug-produced syndromes, meaning the presentation to the emergency department, the resuscitation and critical-care escalation and the disposition decision, belongs to the Perinatal/women's mental health and emergency psychiatry notes.
- The procedural technique of the device-based treatments, meaning how each is administered and the service arrangements it needs, is set out in the ECT and neurostimulation notes. This chapter keeps the modality and its utility across disorders.

IN INDIA

Two of the six second variables carry a decision that is made differently on Indian ground, and both are worth retrieving as decisions rather than as facts. Where the section that owns the testing framework indicates a pretreatment genetic test before a specific anticonvulsant in people of the ancestries at raised risk, the decision is to test before prescribing rather than to react after a severe cutaneous reaction has occurred. Where the patient's own slow-metaboliser phenotype has been determined, the decision is to follow the gene-drug recommendation that result carries. Where only the population frequency of the variant is known, the decision is ordinary cautious prescribing, opening at the lower end of the range and titrating against response, without labelling this patient a slow metaboliser: a frequency in a population is a prior probability and not a result in the person in front of you. In the first of these the second variable is known before the first prescription, and that is where anticipation is at its most complete.

GOOD TO REMEMBER

Finish every prescribing decision in a complicated patient with a spoken sentence: this agent, at this exposure, in the presence of this second variable, watched with this measurement. Said aloud on a ward round it exposes the limb nobody thought through, and it is the sentence that earns marks in a viva. A decision that cannot be completed in that form has not been made pharmacologically.

Carried into the examination, the whole chapter collapses to a sequence that survives when detail fails. Name the second variable the stem introduced. Say whether it changes exposure or effect. State what you would measure because of that, what you would do with the answer, and what you are handing to somebody else. Consolidation buys not more facts but faster access to the ones

already learned, so that under pressure the reasoning arrives as a method rather than as half-remembered warnings about individual drugs.

> Safe psychopharmacology is the management of second variables: name what else is acting on the prescription, decide whether it changes exposure or effect, measure the thing that moves where a measurement will inform the decision, and act on the established risk or the observed change rather than on the label.

References

Every specific in this chapter was taken from a named source, and the figure against each entry below is the number of specifics that came from that work. This section lists the 138 works those 662 specifics rest on. It is a list of works, not of page references: for each specific, the chapter's source record keeps the exact page, table row, source file and date of retrieval, 605 separate entries in all, and that detail is kept there rather than reprinted here. To check where one specific came from, look it up there; to find a work, read this list. Doses were read at current regulatory labelling, and each is filed with the regulator, the product, the label version and the date the label was read. 3 specifics are recorded at partial confidence. This list records where each specific came from. It does not assert that the chapter's prose uses each source correctly, which is checked separately.

The Maudsley Prescribing Guidelines in Psychiatry, 15th edition, 2025
132 specifics in this chapter.

Indian Psychiatric Society Clinical Practice Guidelines 2025. Indian J Psychiatry 2026;68(1). Which disorder-specific guideline and which table each specific came from is recorded in the chapter's source record.
88 specifics in this chapter.

Kaplan and Sadock's Comprehensive Textbook of Psychiatry, 2024 edition
23 specifics in this chapter.

Annual Performance Report of PvPI 2023-2024
22 specifics in this chapter.

Essentials of Medical Pharmacology, K. D. Tripathi
17 specifics in this chapter.

The Maudsley Prescribing Guidelines in Psychiatry, 14th edition, 2021
13 specifics in this chapter.

DSM-5-TR, Fifth Edition Text Revision, 2022
12 specifics in this chapter.

The Maudsley Prescribing Guidelines for Mental Health Conditions in Physical Illness, 1st edition
8 specifics in this chapter.

The Maudsley Prescribing Guidelines in Psychiatry. Edition not stated in these citations; the copy they cite is the one that other citations in this chapter name as the fourteenth edition, Wiley, 2021.
7 specifics in this chapter.

Clinical Methods in Psychiatry, 1st edition, 2024
6 specifics in this chapter.

Spravato (esketamine) nasal spray SmPC, electronic medicines compendium product 10977

5 specifics in this chapter.

The Lithium Handbook, Meyer and Stahl

5 specifics in this chapter.

Goodman and Gilman's The Pharmacological Basis of Therapeutics

4 specifics in this chapter.

Stahl's The Clozapine Handbook, Meyer and Stahl

4 specifics in this chapter.

The Maudsley Guidelines on Advanced Prescribing in Psychosis, 1st edition

4 specifics in this chapter.

BAP guideline 2025

3 specifics in this chapter.

Geriatric Psychiatry Study Guide

2 specifics in this chapter.

Child and Adolescent Psychiatry, Indian text, Malhotra (editor)

1 specific in this chapter.

Malaysian Clinical Practice Guideline on Management of Schizophrenia, 2011

1 specific in this chapter.

PRACTIN (cyproheptadine) tablet prescribing information

1 specific in this chapter.

Individual studies, guidance documents and product labels cited, in alphabetical order. A figure in parentheses is the number of specifics that source carried.

1. Ables AZ, Nagubilli R. Am Fam Physician. 2010;81(9):1139-1142. https://www.aafp.org/pubs/afp/issues/2010/0501/p1139.html (13)
2. Aggarwal R et al. HLA-B*1502 is associated with carbamazepine induced Stevens-Johnson syndrome in North Indian population. Hum Immunol 2014;75(11):1120-2. PMID 25305458, doi:10.1016/j.humimm.2014.09.022 (abstract) (1)
3. Allen MH, Hirschfeld RM, Wozniak PJ, Baker JD, Bowden CL. Linear relationship of valproate serum concentration to response and optimal serum levels for acute mania. Am J Psychiatry 2006;163(2):272-5. DOI 10.1176/appi.ajp.163.2.272 (PMID 16449481, via PubMed) (1)
4. Amstutz U et al. Epilepsia 2014;55(4):496-506. PMID 24597466, doi:10.1111/epi.12564 (abstract, Results) (2)
5. Arumugham SS, et al. IPS Clinical Practice Guidelines for OCD. Indian J Psychiatry 2026;68(1), 'OCD during pregnancy and lactation', point 6 (1)
6. Barik AK, Chauhan R, Reddy A, Kaloria N. Indian J Psychol Med 2024;46(1):94-95. https://doi.org/10.1177/02537176231201660 (PMC10958069) (2)
7. Benet LZ, Hoener BA. Changes in plasma protein binding have little clinical relevance. Clin Pharmacol Ther 2002;71(3):115-21. PMID 11907485. https://doi.org/10.1067/mcp.2002.121829 - TITLE-LEVEL claim only; PubMed carries no abstract for this article and no copy of the full text was available (1)
8. Blumberger DM et al. THREE-D trial. Lancet 2018;391(10131):1683-1692. https://pubmed.ncbi.nlm.nih.gov/29726344/ (1)
9. Blumberger DM, Vila-Rodriguez F, Thorpe KE, et al. Effectiveness of theta burst versus high-frequency repetitive transcranial magnetic stimulation in patients with depression (THREE-D): a randomised non-inferiority trial. Lancet 2018;391(10131):1683-1692. NCT01887782. PMID 29726344 (1)
10. Boyer EW, Shannon M. The serotonin syndrome. N Engl J Med 2005;352(11):1112-20. PMID 15784664. https://doi.org/10.1056/NEJMra041867 - quoted from the (1)
11. Brown JT et al. CPIC Guideline for CYP2D6 Genotype and Atomoxetine Therapy. Clin Pharmacol Ther 2019;106(1):94-102. PMID 30801677, PMC6612570. https://doi.org/10.1002/cpt.1409 - quoted from the (1)

12. Brown S, Rittenbach K, Cheung S, McKean G, MacMaster FP, Clement F. Current and Common Definitions of Treatment-Resistant Depression: Findings from a Systematic Review and Qualitative Interviews. Can J Psychiatry 2019;64(6):380-387 (1)
13. Brunoni AR, Moffa AH, Sampaio-Júnior B, et al. Trial of electrical direct-current therapy versus escitalopram for depression (ELECT-TDCS). N Engl J Med 2017;376(26):2523-2533. PMID 28657871. https://pubmed.ncbi.nlm.nih.gov/28657871/ (1)
14. Bschor T, Kern H, Henssler J, Baethge C. Switching the Antidepressant After Nonresponse in Adults With Major Depression: A Systematic Literature Search and Meta-Analysis. J Clin Psychiatry 2018;79(1):16r10749. https://doi.org/10.4088/JCP.16r10749 (1)
15. Calle AM et al, World Allergy Organ J 2023;16(3):100673 (9)
16. CDSCO, Classification of Medical Devices, Neurological category, serial 33. Document circulated under File No.29/Misc/03/2020-DC (200), Notice dated 13 September 2020 on classification of non-notified medical devices, annexed 'for finalization' with 30 days for stakeholder comments; the table column is headed 'Current CDSCO Classification'. https://cdsco.gov.in/opencms/resources/UploadCDSCOWeb/2018/UploadPublic_NoticesFiles/Final%20Classification%20of%20MD.pdf (1)
17. Chen P et al. N Engl J Med 2011;364(12):1126-33. PMID 21428768, doi:10.1056/NEJMoa1009717 (abstract, Methods) (2)
18. Cole EJ et al. Am J Psychiatry 2022;179(2):132-141. https://pubmed.ncbi.nlm.nih.gov/34711062/ (2)
19. CORRECT-BD pilot trial. https://pubmed.ncbi.nlm.nih.gov/42087328/ (1)
20. CPIC 2018, https://pmc.ncbi.nlm.nih.gov/articles/PMC5847474/ (26)
21. CPIC carbamazepine recommendation table, 'Comments' column, HLA-B*15:02 negative rows. https://files.cpicpgx.org/data/report/current/recommendation/carbamazepine%20recommendation.xlsx (3)
22. CPIC citalopram recommendation table, 'Comments' column, CYP2C19 Poor Metabolizer row. https://files.cpicpgx.org/data/report/current/recommendation/citalopram%20recommendation.xlsx (2)
23. CPIC CYP2C19 Frequency Table, 'Phenotype frequency' sheet. https://files.cpicpgx.org/data/report/current/frequency/CYP2C19_frequency_table.xlsx (linked from https://cpicpgx.org/) (2)
24. CPIC CYP2D6 Frequency Table, 'Methods and Caveats' sheet. https://files.cpicpgx.org/data/report/current/frequency/CYP2D6_frequency_table.xlsx (2)
25. CPIC HLA-B Frequency Table, 'Allele frequency' sheet, *15:02 row. https://files.cpicpgx.org/data/report/current/frequency/HLA-B_frequency_table.xlsx (1)
26. CPIC paroxetine recommendation table, 'CYP2D6 Activity Score' and 'CYP2D6 Phenotype' columns. https://files.cpicpgx.org/data/report/current/recommendation/paroxetine%20recommendation.xlsx (4)
27. CPIC phenytoin 2020 update, 'Other considerations, HLA'. https://pmc.ncbi.nlm.nih.gov/articles/PMC7831382/ (6)
28. CREST-MST trial, NCT03191058. https://pubmed.ncbi.nlm.nih.gov/41997695/ (1)
29. Daly EJ, Trivedi MH, Janik A, et al. Efficacy of Esketamine Nasal Spray Plus Oral Antidepressant Treatment for Relapse Prevention in Patients With Treatment-Resistant Depression: A Randomized Clinical Trial. JAMA Psychiatry 2019;76(9):893-903. https://doi.org/10.1001/jamapsychiatry.2019.1189 (PMC6551577) (1)
30. de Filippis R et al. Screening the European pharmacovigilance database for reports of clozapine-related DRESS syndrome: 47 novel cases. Eur Neuropsychopharmacol 2022;60:25-37. PMID 35635994, doi:10.1016/j.euroneuro.2022.04.009 (abstract) (1)
31. de Filippis R et al. Unravelling cases of clozapine-related DRESS in patients reported otherwise: A systematic review. J Psychopharmacol 2021;35(9):1062-1073. PMID 34044659, doi:10.1177/02698811211021587 (abstract, Results) (1)
32. Deng ZD et al. Neuropsychopharmacology 2024;49(1):150-162. https://pmc.ncbi.nlm.nih.gov/articles/PMC10700353/ (2)
33. Devi K. The association of HLA B*15:02 allele and Stevens-Johnson syndrome/toxic epidermal necrolysis induced by aromatic anticonvulsant drugs in a South Indian population. Int J Dermatol 2018;57(1):70-73. PMID 29076187, doi:10.1111/ijd.13812 (abstract, Conclusions) (1)
34. Duma A, Maleczek M, Panjikaran B, Herkner H, Karrison T, Nagele P. Major Adverse Cardiac Events and Mortality Associated with Electroconvulsive Therapy: A Systematic Review and Meta-analysis. Anesthesiology 2019;130(1):83-91, Abstract (Results), first two sentences quoted contiguously. PMID 30557212. PMC6300062. https://pmc.ncbi.nlm.nih.gov/articles/PMC6300062/ (1)

35. Dunkley EJC, Isbister GK, Sibbritt D, Dawson AH, Whyte IM. QJM. 2003;96(9):635-42. PMID 12925718, https://doi.org/10.1093/qjmed/hcg109 (4)

36. Faber MS, Fuhr U. Clin Pharmacol Ther 2004;76(2):178-84. PMID 15289794. https://doi.org/10.1016/j.clpt.2004.04.003 (3)

37. FDA-approved atomoxetine (Strattera) label, as reproduced in Medical Genetics Summaries, NCBI Bookshelf NBK315951. https://www.ncbi.nlm.nih.gov/books/NBK315951/ (2)

38. FDA-approved label, Lamictal (lamotrigine), GlaxoSmithKline, BOXED WARNING 'SERIOUS SKIN RASHES'. DailyMed setid d7e3572d-56fe-4727-2bb4-013ccca22678, https://dailymed.nlm.nih.gov/dailymed/drugInfo.cfm?setid=d7e3572d-56fe-4727-2bb4-013ccca22678 (4)

39. FDA-approved label, Tegretol (carbamazepine), WARNINGS. https://dailymed.nlm.nih.gov/dailymed/drugInfo.cfm?setid=8d409411-aa9f-4f3a-a52c-fbcb0c3ec053 (11)

40. FDA-approved label, Trileptal (oxcarbazepine), Novartis, section 5.4 'Serious Dermatological Reactions', 'Association With HLA-B*1502'. DailyMed setid 4c5c86c8-ab7f-4fcf-bc1b-5a0b1fd0691b, https://dailymed.nlm.nih.gov/dailymed/drugInfo.cfm?setid=4c5c86c8-ab7f-4fcf-bc1b-5a0b1fd0691b (4)

41. Fekadu A, Wooderson S, Donaldson C, et al. A multidimensional tool to quantify treatment resistance in depression: the Maudsley staging method. J Clin Psychiatry 2009;70(2):177-184. https://doi.org/10.4088/jcp.08m04309 (PMID 19192471) (1)

42. Finley PR, Warner MD, Peabody CA. Clin Pharmacokinet 1995;29(3):172-91, abstract. PMID 8521679. https://doi.org/10.2165/00003088-199529030-00004 (2)

43. Gautam S et al. Indian J Psychiatry 2017;59(Suppl 1):S34-S50. https://pmc.ncbi.nlm.nih.gov/articles/PMC5310101/ (5)

44. Groppa S, Oliviero A, Eisen A, Quartarone A, Cohen LG, Mall V, et al. A practical guide to diagnostic transcranial magnetic stimulation: report of an IFCN committee. Clin Neurophysiol 2012;123(5):858-882. doi 10.1016/j.clinph.2012.01.010. PMID 22349304. PMC4890546. https://pmc.ncbi.nlm.nih.gov/articles/PMC4890546/ (1)

45. Grover S, Chakrabarti S, Kulhara P, Avasthi A. Clinical Practice Guidelines for Management of Schizophrenia. Indian J Psychiatry 2017;59(Suppl 1). https://pmc.ncbi.nlm.nih.gov/articles/PMC5310098/ (2)

46. Gurrera RJ, Caroff SN, Cohen A, et al. J Clin Psychiatry. 2011;72(9):1222-8. PMID 21733489, https://doi.org/10.4088/JCP.10m06438 (2)

47. Gurrera RJ, Mortillaro G, Velamoor V, Caroff SN. A Validation Study of the International Consensus Diagnostic Criteria for Neuroleptic Malignant Syndrome. J Clin Psychopharmacol. 2017;37(1):67-71. PMID 28027111, https://doi.org/10.1097/JCP.0000000000000640 (1)

48. Henssler J, Alexander D, Schwarzer G, Bschor T, Baethge C. Combining Antidepressants vs Antidepressant Monotherapy for Treatment of Patients With Acute Depression: A Systematic Review and Meta-analysis. JAMA Psychiatry 2022;79(4):300-312. https://doi.org/10.1001/jamapsychiatry.2021.4313 (PMC8851370) (1)

49. IPS Clinical Practice Guideline, bipolar disorder, 2025 (Indian J Psychiatry 2026;68) (1)

50. IPS Clinical Practice Guideline, schizophrenia, 2025 (1)

51. IPS Clinical Practice Guidelines for the Use of Electroconvulsive Therapy. Indian J Psychiatry 2023;65(2):258-269. https://pmc.ncbi.nlm.nih.gov/articles/PMC10096214/ (2)

52. Kajumba MM et al. Mol Biomed 2024;5:43. https://doi.org/10.1186/s43556-024-00205-y (PMC11485009) (3)

53. Kaniwa N et al. Epilepsia 2010;51(12):2461-5. PMID 21204807, doi:10.1111/j.1528-1167.2010.02766.x (abstract; quote begins mid-sentence, eliding only the opening discourse marker 'In the current study,') (1)

54. Kennedy SH et al. CANMAT 2016 Clinical Guidelines for the Management of Adults with Major Depressive Disorder: Section 3. Pharmacological Treatments. Can J Psychiatry 2016;61(9):540-560. https://doi.org/10.1177/0706743716659417 (PMC4994790). Prose sentence from section 3.18 'Other medications'; dose from Table 11 (recommendations for adjunctive agents) with footnote b. Both strings verified in the full text 2026-07-31. Bracketed labels in the quote are editorial locators not source text (1)

55. Khor AH, Lim KS, Tan CT, Wong SM, Ng CC. HLA-B*15:02 association with carbamazepine-induced Stevens-Johnson syndrome and toxic epidermal necrolysis in an Indian population: a pooled-data analysis and meta-analysis. Epilepsia 2014;55(11):e120-4. doi:10.1111/epi.12802 (PMID 25266342) (1)

56. Khosravi M. Deep brain stimulation in treatment-resistant psychiatric disorders: efficacy, safety, and future directions. Brain Sci 2025;15(11):1244. doi 10.3390/brainsci15111244. PMID 41300250. PMC12650556 (1)

57. Krishna MBN, et al. Asia Pac Psychiatry 2024;16(2):e12555. DOI 10.1111/appy.12555 (PMID 38565831, via PubMed) (2)

58. Lam RW et al. CANMAT 2023 Update, MDD in Adults. Can J Psychiatry 2024;69(9):641-687. https://doi.org/10.1177/07067437241245384 (PMC11351064), Q.7.g Adjunctive medications (4)

59. Lamberg TS, Kivisto KT, Neuvonen PJ. Br J Clin Pharmacol 1998;45(4):381-5, discussion. PMID 9578186. https://doi.org/10.1046/j.1365-2125.1998.t01-1-00698.x (2)

60. Lilja JJ et al. Clin Pharmacol Ther 1998;64(6):655-60, conclusion. PMID 9871430. https://doi.org/10.1016/S0009-9236(98)90056-X (2)

61. LiverTox, 'Adverse Drug Reaction Probability Scale (Naranjo)'. https://www.ncbi.nlm.nih.gov/books/NBK548069/ (3)

62. Low-intensity tES safety and guidance paper. https://pmc.ncbi.nlm.nih.gov/articles/PMC7618810/ (4)

63. Lu ML, Lane HY, Chen KP, Jann MW, Su MH, Chang WH. Fluvoxamine reduces the clozapine dosage needed in refractory schizophrenic patients. J Clin Psychiatry 2000;61(8):594-9. PMID 10982203. https://doi.org/10.4088/jcp.v61n0809 (1)

64. Luchini F, Medda P, Mariani MG, Mauri M, Toni C, Perugi G. Electroconvulsive therapy in catatonic patients: efficacy and predictors of response. World J Psychiatry 2015;5(2):182-192. PMID 26110120. https://pmc.ncbi.nlm.nih.gov/articles/PMC4473490/ (1)

65. Malhi GS, Bell E, Bassett D, et al. The 2020 Royal Australian and New Zealand College of Psychiatrists clinical practice guidelines for mood disorders. Aust N Z J Psychiatry 2021;55(1):7-117, 'Recommendation Box 1. Physical treatments', item 1.6, grade CBR (1)

66. Mathiassen AB, Semkovska M, Lundsgaard CC, Gbyl K, Videbech P. Autobiographical memory after electroconvulsive therapy: systematic review and meta-analysis. Br J Psychiatry 2025;228(3):263-273. doi 10.1192/bjp.2025.2. PMID 40357797. PMC12916233 (1)

67. Mazzoleni A, Bhatia S, Bantounou MA, Kumar NS, Dzalto M, Soiza RL. Clinical practice guidelines on the use of deep brain stimulation for the treatment of obsessive-compulsive disorder: systematic review. BJPsych Open 2023;9(5):e148. doi 10.1192/bjo.2023.539. PMID 37551586. PMC10486236 (1)

68. McIntyre RS et al. World Psychiatry 2023;22(3):394-412. https://doi.org/10.1002/wps.21120 (3)

69. Mehta TY et al. Association of HLA-B*1502 allele and carbamazepine-induced Stevens-Johnson syndrome among Indians. Indian J Dermatol Venereol Leprol 2009;75(6):579-82. PMID 19915237, doi:10.4103/0378-6323.57718 (abstract) (1)

70. Mental Healthcare Act 2017 (2)

71. Mental Healthcare Act 2017 (Act 10 of 2017) (1)

72. MTA 24-month follow-up data as reported in the child/adolescent pharmacology chapter of a psychiatry reference text (reference 22 in the source chapter) (1)

73. NICE CG185, Bipolar disorder: assessment and management [2014]. https://www.nice.org.uk/guidance/cg185/chapter/Recommendations (5)

74. NICE HealthTech guidance HTG396, recommendation 1.1 (guidance from IPG542, published 16 December 2015, moved unchanged). https://www.nice.org.uk/guidance/htg396/chapter/1-Recommendations (1)

75. NICE IPG530 (published 26 August 2015), recommendation 1.1; moved unchanged to HealthTech guidance HTG382. https://www.nice.org.uk/guidance/htg382/chapter/1-Recommendations (1)

76. NICE IPG679 recommendation 1.1, moved unchanged to HTG551. https://www.nice.org.uk/guidance/htg551/chapter/1-Recommendations ; VA/DoD MDD CPG 2022 Recommendation 33 (1)

77. NICE NG222, Depression in adults: treatment and management (2022). https://www.nice.org.uk/guidance/ng222/chapter/recommendations (4)

78. NICE TA59, Guidance on the use of electroconvulsive therapy. https://www.nice.org.uk/guidance/ta59/chapter/1-Guidance (2)

79. Paribello P et al. Front Pharmacol 2026. https://doi.org/10.3389/fphar.2026.1819237 (PMC13299096) (3)

80. Petrides G et al. Am J Psychiatry 2015;172(1):52-58. https://pubmed.ncbi.nlm.nih.gov/25157964/ (1)

81. Praharaj SK. Indian J Psychiatry 2025;67(10):1007-1012. https://pmc.ncbi.nlm.nih.gov/articles/PMC12594069/ (5)

82. Prakash S, Patel V, Kakked S, Patel I, Yadav R. Ann Indian Acad Neurol. 2015;18(2):226-30, abstract. PMID 26019424, PMC4445202, https://doi.org/10.4103/0972-2327.150612 (2)

83. Ragheb M. J Clin Psychopharmacol 1990;10(5):350-4, abstract. PMID 2258452 (2)

84. Reif A, Bitter I, Buyze J, et al. Esketamine Nasal Spray versus Quetiapine for Treatment-Resistant Depression. N Engl J Med 2023;389(14):1298-1309. https://doi.org/10.1056/NEJMoa2304145 (1)

85. Review of neuromodulation in psychiatry, describing RECOVER. https://pmc.ncbi.nlm.nih.gov/articles/PMC13149262/ (3)

86. Rose D, Fleischmann P, Wykes T, Leese M, Bindman J. Patients' perspectives on electroconvulsive therapy: systematic review. BMJ 2003;326(7403):1363. PMID 12816822. PMC162130. https://pmc.ncbi.nlm.nih.gov/articles/PMC162130/ (1)

87. Rossi S et al., IFCN safety consensus update, Clin Neurophysiol 2021;132(1):269-306. https://pmc.ncbi.nlm.nih.gov/articles/PMC9094636/ (5)

88. Rush AJ, Trivedi MH, Wisniewski SR, et al. Acute and longer-term outcomes in depressed outpatients requiring one or several treatment steps: a STAR*D report. Am J Psychiatry 2006;163(11):1905-1917. https://doi.org/10.1176/ajp.2006.163.11.1905 (1)

89. Rush AJ, Trivedi MH, Wisniewski SR, et al. Bupropion-SR, sertraline, or venlafaxine-XR after failure of SSRIs for depression. N Engl J Med 2006;354(12):1231-1242. https://doi.org/10.1056/NEJMoa052963 (1)

90. Rybak YE et al. Depress Anxiety 2021;38(4):456-467. https://doi.org/10.1002/da.23135 (PMC8049072) (4)

91. Sasidharanpillai S et al. Indian Dermatol Online J 2022;13(1):40-45 (7)

92. Semkovska M, McLoughlin DM. Biol Psychiatry 2010;68(6):568-577. https://pubmed.ncbi.nlm.nih.gov/20673880/ (1)

93. Shah R, et al. Clinical Practice Guidelines for Assessment and Management of Attention-Deficit/Hyperactivity Disorder. In: IPS Clinical Practice Guidelines for Child and Adolescent Psychiatric Disorders. Indian J Psychiatry 2019;61(Suppl 2). Side-effect monitoring table, 'Tachycardia and rise in blood pressure' row (1)

94. Shi Y, Cai G, Wu W. A panoramic review of transcranial focused ultrasound neuromodulation: from basic research to clinical applications. J Neuroeng Rehabil 2025;22(1):227. doi 10.1186/s12984-025-01753-2. PMID 41152970. PMC12570825. Section 'Stimulation parameters and protocols of tFUS' (1)

95. Sinclair DJM, Zhao S, Qi F, Nyakyoma K, Kwong JSW, Adams CE. Electroconvulsive therapy for treatment-resistant schizophrenia. Cochrane Database Syst Rev 2019;3(3):CD011847. PMID 30888709. https://pubmed.ncbi.nlm.nih.gov/30888709/ (1)

96. Singh A, Kar SK. How Electroconvulsive Therapy Works? Clin Psychopharmacol Neurosci 2017;15(3):210-221. PMC5565084. https://pmc.ncbi.nlm.nih.gov/articles/PMC5565084/ (1)

97. Singh S et al. A prospective observational study validating the CET score as a screening tool in suspected DRESS syndrome. Int J Dermatol 2024;63(9):1178-1184. PMID 38415838, doi:10.1111/ijd.17080 (abstract) (1)

98. Soldini A, Vogelmann U, Aust S, et al. Neurocognitive function as outcome and predictor for prefrontal transcranial direct current stimulation in major depressive disorder: an analysis from the DepressionDC trial. Eur Arch Psychiatry Clin Neurosci 2025;275(6):1715-1724. doi 10.1007/s00406-024-01759-2. PMID 38407625. PMC12500756 (1)

99. Stahl SM. Stahl's Prescriber's Guide, 6th edition, memantine drug entry #72, pp. 435-438, as recorded in a workspace clinical entry whose Ra-authored wording is the entry's own and is not printed Stahl text (1)

100. Stahl SM. Stahl's Prescriber's Guide, 6th edition, rivastigmine drug entry #111, pp. 655-659, as recorded in a workspace clinical entry whose Ra-authored wording is the entry's own and is not printed Stahl text (1)

101. StatPearls, Anticholinergic Toxicity, NBK534798. https://www.ncbi.nlm.nih.gov/books/NBK534798/ (6)

102. StatPearls, Serotonin Syndrome, NBK482377. https://www.ncbi.nlm.nih.gov/books/NBK482377/ (4)

103. Suhas S, Kumar V, Damodharan D, et al. Do Indian patients with schizophrenia need half the recommended clozapine dose to achieve therapeutic serum level? An exploratory study. Schizophr Res 2020;222:195-201. doi:10.1016/j.schres.2020.05.057 (PMID 32518001) (1)

104. Tikka SK et al. Indian J Psychiatry 2023;65(2):270-288 (17)

105. Trivedi MH, Fava M, Wisniewski SR, et al. Medication augmentation after the failure of SSRIs for depression. N Engl J Med 2006;354(12):1243-1252. https://doi.org/10.1056/NEJMoa052964 (1)

106. Tse L et al. Curr Neuropharmacol. 2015;13(3):395-406, PMC4812801 (9)

107. UK ECT Review Group. Lancet 2003;361(9360):799-808. https://pubmed.ncbi.nlm.nih.gov/12642045/ (2)

108. Uppsala Monitoring Centre, 'About VigiBase'. https://who-umc.org/vigibase-data-access/about-vigibase/ (1)

109. Uppsala Monitoring Centre, 'What is a signal?'. https://who-umc.org/signal-management/what-is-a-signal/ (2)

110. US FDA Humanitarian Device Exemption H050003 (Medtronic). https://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfHDE/hde.cfm?id=H050003 (1)

111. US FDA PMA P960009/S478 (Medtronic). https://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfpma/pma.cfm?id=P960009S478 (1)

112. US FDA PMA P970003/S050 (LivaNova/Cyberonics). https://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfpma/pma.cfm?id=P970003S050 (1)

113. US FDA Summary of Safety and Effectiveness Data, P970003/S050. https://www.accessdata.fda.gov/cdrh_docs/pdf/P970003S050B.pdf (1)

114. US FDA, Drug Development and Drug Interactions Table of Substrates, Inhibitors and Inducers. https://www.fda.gov/drugs/drug-interactions-labeling/drug-development-and-drug-interactions-table-substrates-inhibitors-and-inducers (6)

115. VA/DoD Clinical Practice Guideline for the Management of Major Depressive Disorder (2022). https://www.healthquality.va.gov/guidelines/MH/mdd/ (1)

116. VA/DoD Clinical Practice Guideline for the Management of Major Depressive Disorder (2022). https://www.healthquality.va.gov/guidelines/MH/mdd/VADODMDDCPGFinal508.pdf (3)

117. WHO/UMC, 'The use of the WHO-UMC system for standardised case causality assessment'. https://www.who.int/docs/default-source/medicines/pharmacovigilance/whocausality-assessment.pdf (5)

118. Yadav AK, Chakkumkollath AK, Helna A, et al. Indian Heart J 2023;75(5):343-346. doi:10.1016/j.ihj.2023.05.005 (PMID 37217095) (2)

PART IX

IX

Recent Advances and Landmark Studies

What changed recently and which studies changed it, read as evidence rather than as news.

01 Recent advances and landmark studies

02 Individual newer agents

03 Consolidate and apply

VISUAL 9.1



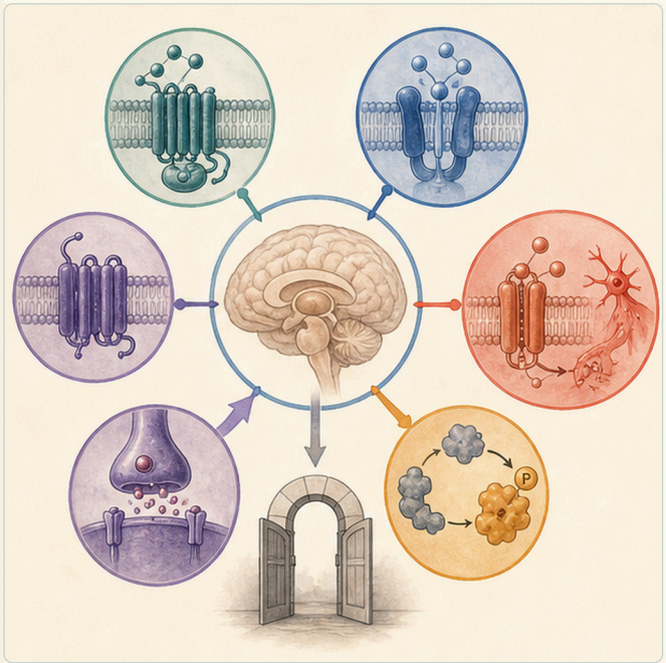
Rapid antidepressants use different treatment containers

- Glutamatergic
- Neurosteroid
- Monoamine
- Setting and monitoring
- Continuation plan

Rapid-acting antidepressant classes differ in mechanism, indication, delivery setting and continuation plan; speed alone does not make them interchangeable.

Source lock: day-38.final.html#sec-t04

VISUAL 9.2



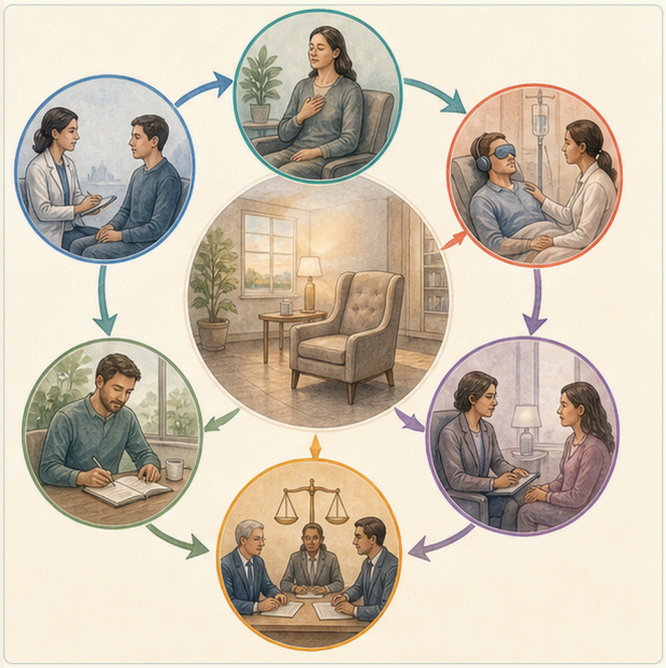
New mechanisms widen psychiatry beyond monoamines

- Muscarinic
- Orexin
- Sigma and NMDA
- Dopamine partial agonism
- PKC targeting
- Approval and indication

Recent agents expand mechanism classes, but each claim remains bounded by the medicine, indication, regulator and evidence base.

Source lock: day-38.final.html#sec-t13

VISUAL 9.3



Psychedelic treatment is more than a molecule

- Screening
- Preparation
- Supervised dosing
- Psychological support
- Integration
- Regulatory status

Psychedelic-assisted therapy is a structured intervention whose screening, support, setting, follow-up and regulatory status are part of the treatment.

Source lock: day-38.final.html#sec-t10

VISUAL 9.4



Computational psychiatry needs an evidence ladder

- Data
- Model
- External validation
- Bias and privacy
- Clinical integration
- Regulatory evidence

AI and computational tools require valid data, external testing, bias and privacy safeguards, and clinical integration before they can improve care.

Source lock: day-38.final.html#sec-t12

PAPER IV

Recent advances and landmark studies

Two bands . one complete notes document

Recent advances and landmark studies

1 · Recent advances and landmark studies for the psychiatrist: what this chapter owns, and what it points to

A drug is never simply approved: it is approved somewhere, for something, on a named authority, as of a date. This chapter gathers the agents, arguments and technologies that arrived in psychiatry most recently. The obvious risk in such a chapter is overlap, since most of these molecules are taught at depth elsewhere in these papers, and the boundaries later on this page handle it. The risk that actually costs marks is different, and installing it is this page's purpose: many of the questions this chapter can ask are about a status rather than a mechanism, and a status is not a property of a molecule.

That is a claim about most of the chapter and not about all of it, and the limit tells you where the discipline applies. Thirteen of the eighteen sections ahead have a regulatory record behind them, at the grain of one agent in one jurisdiction. Five do not, and not through oversight: an argument about a body of evidence, three technologies that are not medicines, and one landmark whose status is left deliberately to the chapter that owns it. In those five the examinable content is mechanism, appraisal and readiness, and a candidate looking for an approval to quote will find none and must not supply one. So the first move on any stem is to decide which of the two kinds of question you have.

1.1 The lens: a status claim is a four-field claim

A statement that an agent may be given, to whom and for what, is a claim about a regulator's document rather than about pharmacology, and this chapter treats it as a **four-field status claim**. The **four fields** are the jurisdiction, the exact indication or trial purpose in the source's own words, the authoritative source that carries it, and the date on which the status was true, which these notes record as **status true as of**. A sentence missing any one of the four is unmarkable, for a clinical reason rather than a stylistic one.

Jurisdiction	Indication	Source	Date
WHICH REGULATOR	IN THE SOURCE'S OWN WORDS	THE DOCUMENT, NAMED	WHEN IT WAS TRUE

Each field fails in a characteristic way. Drop the jurisdiction and the sentence is true somewhere and false where the patient is. Paraphrase the indication and the population quietly widens, because a licence is written for a defined group and the defining words are the ones that were reviewed. Leave the source unnamed and the claim rests on the speaker, which under examination means on memory. Leave the date off and a regulatory record is offered as a

permanent fact, when the reason this chapter exists is that these records moved recently and will move again.

RULE **An approval verb never travels alone.** Before writing that an agent is approved, indicated, authorised or licensed, put all four fields beside it: the jurisdiction, the indication in the source's own words, the document, and the date it was read. If any of the four is unavailable the sentence is not written shorter; it is written as a declared gap.

1.2 Jurisdiction is the hazard, and it does not run as a ladder

The organising fact of this chapter is that its agents are carried in up to **three jurisdictions** and that the cell, not the agent, is the unit of truth. Nothing about the three is ordered. The ledger behind these notes carries a grounded Indian row for lemborexant while its European cell is a declared gap on the centralised register, and a grounded Indian row for endoxifen while its United States cell is a declared gap after multi-route probing of the federal application database. Reading the United States first, the European Union next and India last as a ladder of maturity produces two confident wrong answers in this chapter alone.

The intuition behind the ladder is not stupid, which is why it survives. Volume does run that way: across the sixteen agents the ledger holds, the United States column carries the most grounded cells, the European Union fewer, India fewest. But a column total says nothing about any particular cell, and **regulatory recency** is the hazard this chapter names because a recent record is distributed unevenly, for reasons with little to do with maturity. An agent may reach one regulator first because that is where its sponsor filed, or reach India by a route that never went to a large market. Reasoning from column to cell is arithmetic where the question asked for a document.

TRAP **Filling an empty cell from a full one.** Asked about Indian availability, candidates answer from the American label, because the American record is the one they have read. The two cells are independently sourced, and this chapter holds pairs that point in opposite directions, so an inference from one to the other is not a cautious estimate. It is a fabricated citation wearing a real document's name. Say what the Indian record shows, or say what was searched and on whose authority a claim would rest.

1.3 Three states a cell can be in

Every cell reads in one of three ways, and knowing which one you have settles what may be written from it. A **declared gap** is the state most often misread, because it is a statement about an instrument and never about the world. Where the search ran on several routes, each positive-controlled in the same run, these notes call the result a **measured absence** rather than a probe silence, and even then one gap in this chapter turned out to be a probe miss once the record was retrieved another way. The third state is the one a candidate never expects: a **supportive source** that recites an approval without being one.

THE CELL READS	WHAT IT LICENSES	WHAT IT NEVER LICENSES
Grounded	The indication in the source's own words, for that jurisdiction alone, with its document and its date beside it	A sentence about any other jurisdiction, or the same sentence later without a fresh reading
Declared gap	What was searched, on which instrument, how sensitively, and whose authority a claim would ultimately rest on	Any categorical negative. Silence on an instrument is a fact about the instrument
Supportive source only	That a named body recited an approval, and the document in which it recited it	That the regulator granted it, or the recited date reused as an approval date

The middle state carries most of the weight, because this chapter has already had one gap turn out to be wrong: a cell that read as an absence was re-probed by another route and the record came back. So a declared gap is written as a description of a search rather than a verdict, on the terms the table above sets. A gap written that way can be closed by the next reader of the primary document. A gap written as a negative has to be retracted.

1.4 The instrument, and the five sections it cannot see

The shape of the apparatus behind these notes decides which questions it can answer. The record holds one row per agent per jurisdiction, sixteen agents across three regulators, every row carrying the same four fields. Those rows attach to thirteen of the chapter's eighteen sections. The five without them are the sections whose subject is not a licensed medicine: the serotonin hypothesis debate, an argument about a body of evidence rather than about any agent's licence; the artificial intelligence, biomarker and digital therapeutics sections, which are not what a drug-approval register records; and the anti-amyloid antibodies, where no row is written deliberately, so that none can be quoted from here.

The instruments behind this chapter are drug-approval instruments. When one returns nothing about a technology that is not a medicine, the silence is a fact about the instrument and not about the product, and a sentence converting it into a status is an invented claim wearing a real register's name. In those five sections the honest move is to say what the instruments cover and what they cannot see, then teach what is genuinely examinable: what the object is, what evidence would establish it, and what a service would have to know before using it.

IN INDIA

The Indian cell is the one most often filled in from memory, and the published CDSCO approval lists behind these notes carry a measured limit worth knowing before any silence in them is trusted: their positive controls all sit in documents dated well before the era this chapter covers, and one component substance of an agent taught here appears nowhere in that corpus under any spelling. So an Indian cell is either written from a document that was read, or declared as a gap naming what was searched and how sensitively. A brand on a pharmacy shelf, a CTRI trial registration, a no-objection certificate for a clinical trial and a committee minute recommending approval are none of them a marketing authorisation, and one agent here sits in exactly that state: a national committee's minutes recite an approval, and the row recording them licenses only the fact that they recite it.

1.5 A map of the chapter

With the lens in place the architecture becomes legible. The chapter runs in two declared halves: the recent advances of the last two years, then six individual newer agents held as whole-agent profiles. The map groups the eighteen sections into seven bands. It is not a list to memorise; it tells you which question a later section answers and which neighbour it borders.

BAND OF THE CHAPTER	THE SECTIONS IN IT	WHAT THE BAND IS FOR
The rapid-acting antidepressants	Esketamine and intranasal therapy; zuranolone and the neurosteroid antidepressants; rapid-acting antidepressants as a class	A comparative class map, so the category is not learned as a third ketamine monograph
Sleep-wake pharmacology	Dual orexin receptor antagonists and sodium oxybate	The sharpest regulatory contrast here: one class, four agents, three patterns across three jurisdictions
Mechanisms outside the established algorithms	Xanomeline-trospium and the muscarinic antipsychotics; endoxifen and the protein-kinase-C-targeted agents	Two mechanisms beside the established algorithms rather than inside them, and the most Indian section here
A landmark named and handed on	Anti-amyloid monoclonal antibodies	Why the class matters to psychiatry and not only to neurology, with its status delegated explicitly
Two arguments about evidence	The serotonin hypothesis debate; psychedelic-assisted therapy	What a body of evidence does and does not license, and how to appraise an intervention that is a package, not a molecule
Research to clinic	Artificial intelligence and machine learning; biomarker-driven and computational psychiatry; digital therapeutics and prescription apps	One readiness ladder read three times: analytical validity, clinical validity, clinical utility, regulatory authorisation, routine deployment
The individual newer agents	Asenapine; lurasidone; brexpiprazole; cariprazine; lumateperone; zuranolone as an agent profile	Whole-agent profiles read cell by cell, with the class primer given once instead of six times

Read the map both ways. Given a question, it says which band holds the answer and what kind of answer is wanted. Given a section, it says what its neighbours already cover, which stops an answer drifting into material being marked elsewhere.

1.6 What this chapter owns, and what it hands over

Naming the owning chapter is faster and better rewarded than rebuilding its material from memory. For the drug handovers below, that chapter holds the clinical pharmacology and this chapter holds only what changed, where, and on whose record. For the last two the neighbour holds the neurobiology and the service question instead, and this chapter holds no regulatory record for either, as the second paragraph above says. The handovers are not the whole picture, and treating them as though they were is its own error: this chapter owns seven of its eighteen sections in full, being the neurosteroid material and its agent profile, the muscarinic, endoxifen and lumateperone sections, and the artificial intelligence and biomarker sections. The research-to-clinic readiness ladder those last two build with the digital therapeutics section is this

chapter's as well. Where a section is owned in full there is no neighbour to point at, and an answer that gestures at another chapter has left the marks on the table.

- Ketamine and esketamine in treatment-resistant depression are taught in full in the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes, and the antidepressants and antipsychotics themselves in the Mood disorders: pharmacotherapy notes.
- The hypnotics and the oxybates belong to the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes and the Eating and sleep-wake disorders notes; the psychedelics as substances belong to the Substance use disorders notes.
- Anti-amyloid monoclonal antibodies are named here and taught in the Dementias and neurocognitive disorders notes, which own their status entirely; this chapter states none.
- Service-level deployment of anything digital belongs to the Community psychiatry and public health notes, and the serotonin hypothesis as neurobiology to the Mood disorders: classification, depressive and bipolar syndromes, neurobiology notes and the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

PITFALL

Quoting a status you learned once as though it were still true. Every row behind this chapter carries the date its status was true, because this is the most recently changed material in the syllabus: an indication is widened, a second regulator authorises the same agent for something narrower, a product is discontinued while its approval stands. A candidate who memorised a status and reproduces it much later has not made a small error of currency: the sentence carries no date, so it offers a stale reading as a permanent fact. Learn the four fields and the shape of each record rather than the record itself, and say when you last read it. In the clinic it means checking the document before prescribing a recently licensed agent, not after a query is raised.

1.7 What the examination rewards

A small number of moves earn most of the marks here, and they recur across every band in different clothing. Naming them before the detail arrives is worth more than any single fact, because the detail in this chapter is unusually perishable and the moves are not.

- Answering at the cell rather than at the agent, so a question about one jurisdiction is never answered out of another's record.
- Producing the four fields as a sentence, instead of a verdict with the qualification left to be assumed.
- Reading a silence as a statement about an instrument, and saying what was searched, how sensitively, and on whose authority a claim would rest.
- Giving an ordering with its basis named wherever no quantity is available, rather than a figure that would have to be invented to be given.
- Deciding whether the stem wants a status or a construct, then handing over what a neighbour owns in one sentence rather than rebuilding it.

Those five are also, in the same order, the five ways an answer here goes wrong, which is why they are habits rather than advice. A candidate who runs them will sometimes have to say that a record was searched and not found, or that an ordering is all the evidence supports. Both are complete answers here, because the alternative is a confident sentence nobody can source.

GOOD TO REMEMBER

Before recommending, prescribing or teaching any agent from this chapter, ask four questions in order. Which regulator's document supports this use in the country where this patient is being treated. What does that document say the indication is, in its own words rather than in yours. When did you last read it. And is the thing you are about to say a status claim, needing all four fields, or a mechanism claim, needing none of them and needing a stated limit instead. A clinician who answers those four before speaking will not be the person who tells a family a treatment is available on the strength of a label written for another country.

1.8 Holding it together

Under examination pressure the detail goes first, because a date and a verbatim indication are harder to hold than a mechanism. What has to survive is a routing sequence. The scaffold below is a teaching device for that sequence rather than a validated instrument, and it will structure an answer on any of the eighteen sections.

MNEMONIC

Kind, Currency, Grain, Owner

Kind: which of the chapter's three kinds of section the stem sits in, an agent with a regulatory record, an argument about a body of evidence, or a technology on a readiness ladder. **Currency**: what the answer is paid in, a status with its four fields, or a construct with its mechanism, its evidence and its limit. **Grain**: if it is a status, answer at one agent in one jurisdiction, and say which of the three states that cell is in. **Owner**: who holds the clinical pharmacology, which for most sections is a neighbouring chapter named in a sentence, and for seven of the eighteen is this chapter, where a full account is owed and a cross-reference earns nothing.

For a short answer, run the four and stop. For a long essay, expand each into a paragraph and take the specifics from the section that owns them. In a viva, make the sequence audible: name the jurisdiction, say what the document says and when it was true, and say what you would have to read to answer for anywhere else. That last sentence separates a candidate who has understood this chapter from one who has memorised it, because it concedes a limit without conceding the question.

Every status sentence carries its jurisdiction, its exact indication, its authoritative source and the date it was true, and no status is ever carried across a border.

2 · Esketamine and intranasal therapy, including monotherapy indication updates

An approval is not a property of a molecule; it is a sentence a named regulator wrote in a named country on a named date. Intranasal esketamine is where that discipline is examined, and the marks are lost in one move: a status carried across a border, almost always from the United States inward, into a sentence true somewhere and false where the patient is. Delivery, dissociation and blood-pressure monitoring, the observation period and the access problem belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes and are not repeated here.

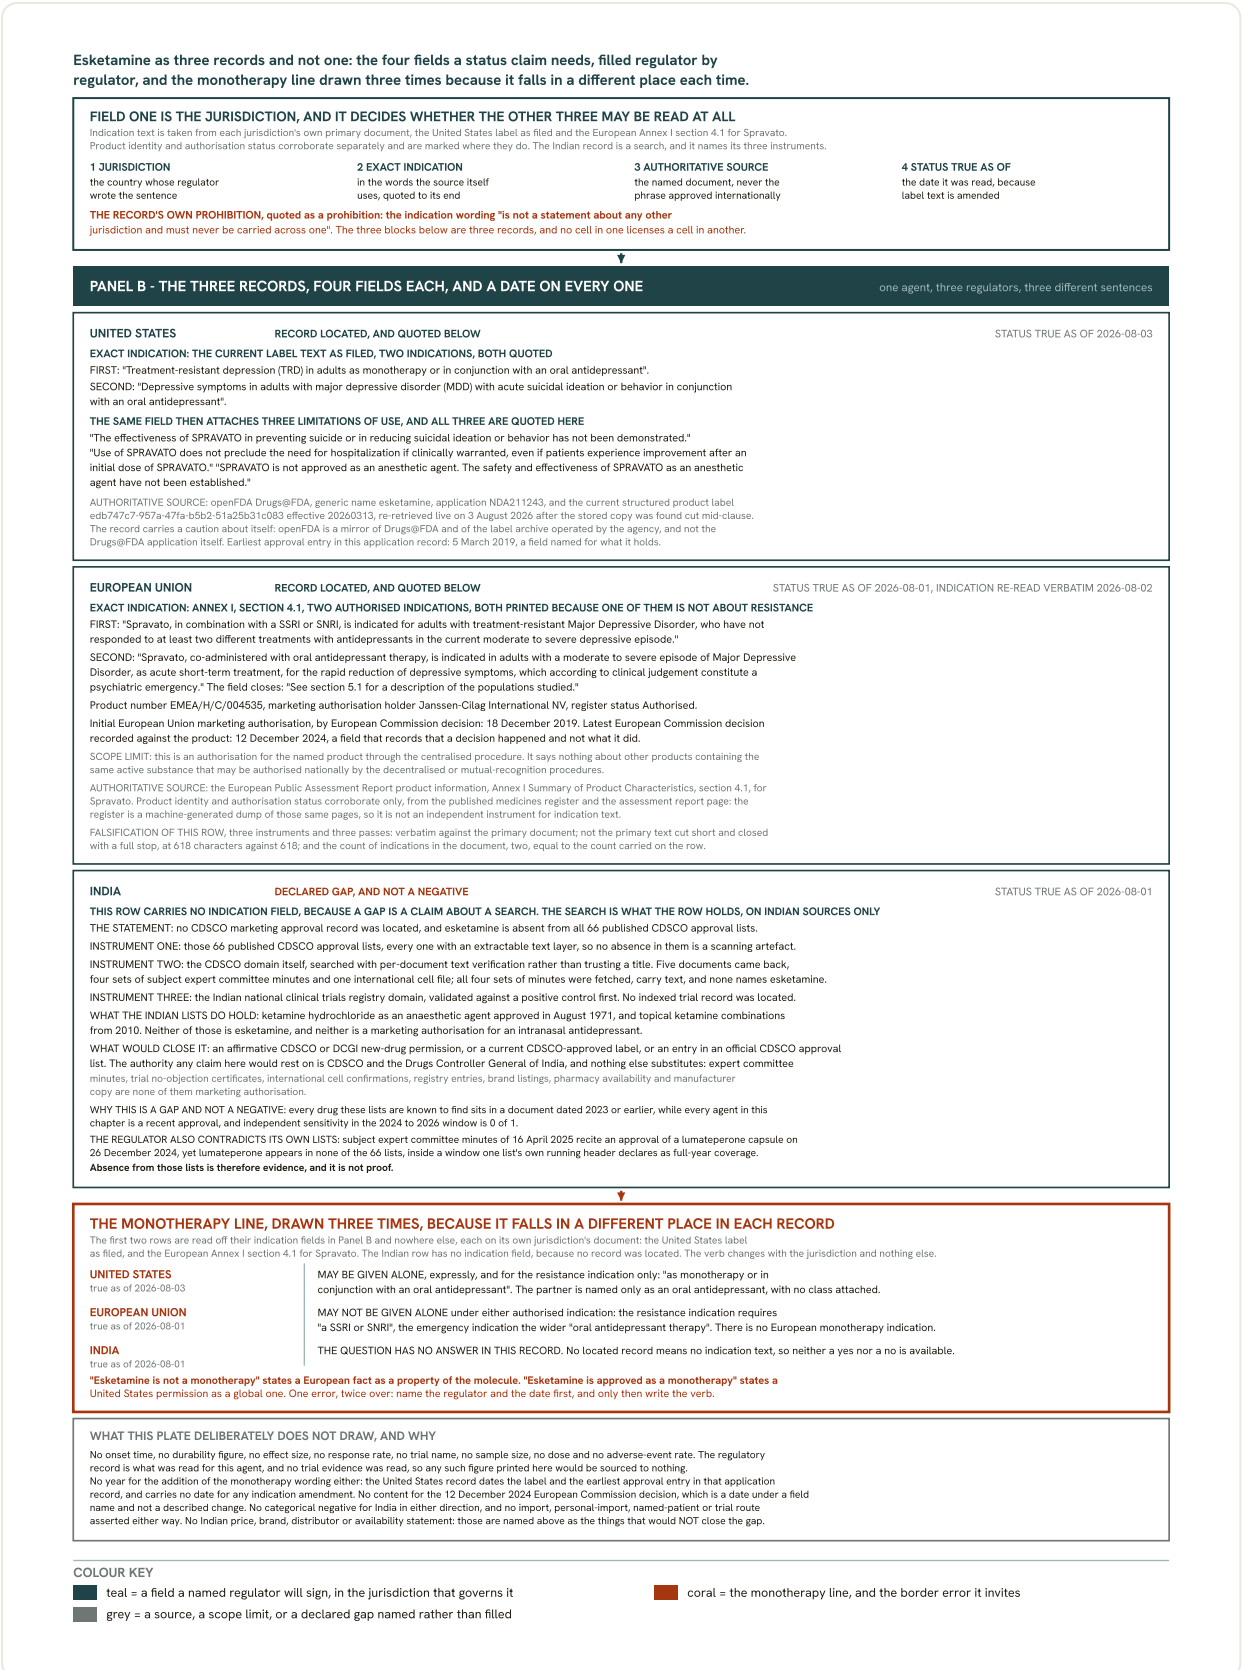


Fig 38.1 **The four-field status ledger applied to esketamine, one record per regulator.** The United States record, read 3 August 2026, permits the spray alone for its resistance indication, as monotherapy or in conjunction with an oral antidepressant. The European record, read 1 August 2026, carries two authorised indications, one requiring a SSRI or SNRI and the other oral antidepressant therapy, so neither authorises it alone. The Indian record is a declared gap: no CDSCO marketing approval record located, across 66 published approval lists, with the document that would close it named. The monotherapy question therefore has three answers and no single one. (openFDA Drugs@FDA NDA211243; the European Annex I section 4.1 for Spravato; the CDSCO approval lists and the Indian trials registry.)

2.1 Four fields, and the sentence that fails without them

A regulatory status is usable only when it carries four fields together: the jurisdiction, the exact indication in the words the source itself uses, the authoritative source, and the date on which the status was true. "Esketamine is approved for depression" fails on all four, and last on the date, in a domain where label text is amended and last year's correct answer can already be wrong.

Every statement below was read from its primary source on 1 August 2026, with the European indication text re-read verbatim on 2 August 2026 and the United States label text re-retrieved on 3 August 2026, and each is offered as true as of that reading and nothing more.

A status never travels, and the source record says so: the label wording is not a statement about any other jurisdiction and must never be carried across one. And "approved internationally" is not an answer: name the regulator, or say that no record was located and what was searched.

RULE **A status sentence is written in one jurisdiction and stays there.** An American indication does not imply a European one, and neither implies anything about a country whose regulator was never asked. Where the record for a jurisdiction is empty, say the record is empty, in that jurisdiction, on that date.

Marketing authorisation is a permission granted by a national regulator, and it is not availability. Expert committee minutes, trial no-objection certificates, registry entries, brand listings and pharmacy availability are none of them marketing authorisation.

2.2 The United States: the indication that carries monotherapy

Two indications sit on the current United States label. The first is for the treatment of treatment-resistant depression (TRD) in adults as monotherapy or in conjunction with an oral antidepressant. The second is for depressive symptoms in adults with major depressive disorder (MDD) with acute suicidal ideation or behavior in conjunction with an oral antidepressant. The resistance indication offers a choice, **monotherapy** or combination; the suicidality indication offers none, the oral antidepressant being part of it rather than an option inside it.

The label then attaches three Limitations of Use. The first: the effectiveness of SPRAVATO in preventing suicide or in reducing suicidal ideation or behavior has not been demonstrated. So the licence is for depressive symptoms in a patient who has **acute suicidal ideation or behavior**, and it is not a licence for treating the suicidality; sliding from the first formulation to the second claims an effect the label explicitly declines. The second: use of SPRAVATO does not preclude the need for hospitalization if clinically warranted, even if patients experience improvement after an initial dose of SPRAVATO. A rapid response is not a disposition decision. The third: SPRAVATO is not approved as an anesthetic agent. The safety and effectiveness of SPRAVATO as an anesthetic agent have not been established.

The status was read from the openFDA Drugs@FDA interface, application NDA211243, with the current structured product label effective 13 March 2026, on 3 August 2026. The record carries a caution about itself: openFDA is a mirror of Drugs@FDA and of the label archive operated by the agency, and not the Drugs@FDA application itself.

The earliest approval entry in that application record is dated **5 March 2019**, and the field is named for what it holds, the earliest approval-status date in the record, not for what a reader would like it to mean. The record carries no date for the monotherapy wording: the label in force at the reading contains it, and nothing in the source read here says when it entered. So the defensible sentence is that the United States indication included monotherapy on the label effective 13 March 2026, read on 3 August 2026; the indefensible one names the year monotherapy was added.

GOOD TO REMEMBER

An indication that permits monotherapy permits it; it does not recommend it, and it does not make the evidence for the drug alone equal to the evidence for the drug alongside an antidepressant. The maintenance trial and the head-to-head comparison taught in the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes are both add-on designs there, so a durability or comparator figure from either carries the design it came from and does not attach to a monotherapy sentence it was never generated under.

2.3 The European Union: two authorised indications, and neither is a monotherapy

The European authorisation covers the same product and is not the same licence. Annex I of the **Summary of Product Characteristics**, section 4.1, carries two therapeutic indications, and both must be reproduced before the licence has been described. The first: Spravato, in combination with a SSRI or SNRI, is indicated for adults with treatment-resistant Major Depressive Disorder, who have not responded to at least two different treatments with antidepressants in the current moderate to severe depressive episode. The second: Spravato, co-administered with oral antidepressant therapy, is indicated in adults with a moderate to severe episode of Major Depressive Disorder, as acute short-term treatment, for the rapid reduction of depressive symptoms, which according to clinical judgement constitute a psychiatric emergency.

Neither indication authorises the spray by itself: the first requires an SSRI or SNRI, the second requires oral antidepressant therapy, so there is no European monotherapy indication of any kind. And the sentence that European esketamine is indicated only for treatment-resistant depression is false, because the second indication is not about resistance. It is an acute, short-term indication for what clinical judgement calls a **psychiatric emergency** inside a moderate to severe episode, and a patient can meet it at first presentation, having failed nothing.

That the second indication is really there was tested. The European row was checked against the primary document on three instruments: whether its text is verbatim in the product information, whether it is the primary text cut short and closed with a full stop so that nothing signals the loss, and whether the count of indications carried matches the count in the document. Two indications are in the document and two are on the row. A one-indication summary of a two-indication licence reads as complete, which is why the third instrument counts.

Spravato holds product number EMEA/H/C/004535, its marketing authorisation holder is Janssen-Cilag International NV, and its status on the European Medicines Agency register is Authorised. The initial European Union marketing authorisation, by European Commission

decision, is dated **18 December 2019**. The most recent European Commission decision recorded against the product is dated **12 December 2024**, and that date needs a warning: the field records the latest decision and does not state what it did. Nothing read here says it touched the indication, and converting a decision date into an indication change manufactures the change out of a timestamp.

One limit belongs on every European sentence here. This is an authorisation for the named product through the **centralised procedure**, and it says nothing about other products containing the same active substance that may be authorised nationally by the decentralised or mutual-recognition procedures, which the central register does not cover by construction. "Authorised in the European Union" therefore means authorised centrally, for this product, and an answer should say so.

2.4 The monotherapy line, drawn regulator by regulator

The line falls differently in each jurisdiction.

JURISDICTION	THE RESISTANCE INDICATION, IN THE SOURCE'S OWN TERMS	MAY THE SPRAY BE GIVEN ALONE?	NAMED PARTNER DRUG
United States	Treatment-resistant depression in adults, as monotherapy or in conjunction with an oral antidepressant	Yes, expressly, for this indication	An oral antidepressant, no class named
European Union	Treatment-resistant major depressive disorder in adults who have not responded to at least two different treatments with antidepressants in the current moderate to severe depressive episode	No	A SSRI or SNRI, named and restricted
India	No CDSCO marketing approval record located as of 1 August 2026	The question has no answer in the record	The question has no answer in the record

One difference is not in the table. Only the European indication writes its own entry criterion, counting failed antidepressant treatments within the current episode; the American resistance indication names the condition and leaves the counting elsewhere.

Both second indications require an oral antidepressant alongside, and there the agreement stops: the American is anchored to a symptom on which the label then declines to claim an effect, the European to a clinical judgement about the situation as a whole.

1 of 2	2	66	0
US INDICATIONS PERMITTING MONOTHERAPY	AUTHORISED EU INDICATIONS, BOTH COMBINATION	CDSCO APPROVAL LISTS SEARCHED	INDIAN APPROVAL RECORDS LOCATED

TRAP **Writing "esketamine is not a monotherapy" as a fact about the drug.** It is a fact about a jurisdiction. The add-on-only sentence matches the European wording, where both authorised indications require a partner antidepressant, and it does not match the current United States label, where the resistance indication reads as monotherapy or in conjunction with an oral antidepressant. It is also the sentence carried in the Mood disorders: pharmacotherapy notes, and that adjunctive position holds on the European wording; the American half of it is what this dated update revises. The mirror-image error is just as costly: writing that esketamine is approved as a monotherapy, full stop, states an American permission as though it were a global one.

TRAP **Writing that the European licence requires an SSRI or SNRI.** True of the resistance indication and untrue of the other one. The SSRI or SNRI restriction is printed in the treatment-resistant indication; the psychiatric-emergency indication says co-administered with oral antidepressant therapy, the wider phrase. A candidate who has learned only the resistance indication will state the narrow partner rule for a patient it was never written for. Quote the indication you mean, and quote it whole.

2.5 India: what was searched, what was not found, and what would close it

The Indian position is the one an Indian examination cares about most and the one most often invented. Stated exactly: no CDSCO marketing approval record was located, and esketamine is absent from all 66 published CDSCO approval lists, as of 1 August 2026. That is a statement about a search, which has to be described before it means anything.

Three instruments were run. The first was the corpus of 66 published CDSCO approval lists, every one of which yielded an extractable text layer, so no absence in it is an artefact of an unreadable scan. The second was a search of the regulator's own domain with each returned document downloaded and its text verified rather than its title trusted. The third was a search of the national clinical trials registry domain, validated against a positive control before any zero result was believed. The design point: a search never shown to find what it should find cannot prove that something is missing.

Five documents came back: four sets of subject expert committee minutes for neurology and psychiatry, and one international cell file. All four sets of minutes were fetched, all carry an extractable text layer, and none of them names esketamine. Those **subject expert committee minutes** were search-index false positives, not evidence. No indexed trial record was located on the national registry either.

The corpus is not empty of the parent molecule, which is where a careless reader manufactures an approval. It contains ketamine hydrochloride as an anaesthetic agent approved in August 1971, and topical ketamine combinations from 2010. Neither of those is esketamine. Neither a racemic anaesthetic licensed for another purpose nor a topical combination licensed for another route is a marketing authorisation for an intranasal antidepressant; and the American label states in terms that SPRAVATO is not approved as an anesthetic agent.

What would close the gap is nameable, which is the test of an honest one. The missing document is an affirmative CDSCO or DCGI new-drug permission, or a current CDSCO-approved label, or

an entry in an official CDSCO approval list. Nothing else substitutes. Expert committee minutes are supportive at best and are not themselves marketing authorisation; a clinical-trial no-objection certificate, a written confirmation from the international cell, a CTRI registration, a brand listing, pharmacy availability and manufacturer copy are none of them marketing authorisation either. Name the document that would settle it, and say whether you have it.

IN INDIA

The sentence to write, in an examination and in a clinical note, is that no CDSCO approval record for intranasal esketamine was located as of August 2026, followed by what was searched: the published approval lists, the regulator's own domain, and the trial registry. Do not write that esketamine is unapproved in India, which is stronger than the evidence supports, and do not write that it is approved, because nothing retrieved says so. A brand name, a price or a distributor's letter is not that document. For a patient or a family: this treatment has a licence elsewhere and no located Indian licence, and what a service can actually offer today is the question to answer first.

2.6 Why the Indian finding is a declared gap and not a negative

A **declared gap** states that a search of named instruments did not retrieve a document, and names the document it wanted. A negative states that the thing does not exist. The first is defensible on the evidence in hand; the second needs an instrument whose silence means something, and that sensitivity has to be measured, not assumed.

Here it was measured, and it is poor in exactly this chapter's window. Every drug the corpus is known to find sits in a document dated 2023 or earlier, while every agent in this chapter is a recent approval, and independent sensitivity in the 2024 to 2026 window is 0 of 1. The corpus is therefore no warrant for a gap on a recently approved agent: it has never been shown to find one.

The stronger leg is that the regulator's own documents contradict its own lists. CDSCO subject expert committee minutes of 16 April 2025 state that a lumateperone capsule had already been approved in the country on 26 December 2024, yet lumateperone appears in none of the 66 lists, and the list that declares full-year 2024 coverage in its own running header carries entries dated 17 and 30 December 2024, so it is not truncated before the approval those minutes recite. A list that misses an approval its own committee recites, inside its own declared coverage window, is a list whose silence cannot be read as absence.

Write the conclusion narrowly. Absence from this corpus is evidence and not proof, and the claim being made is that no CDSCO approval record was located, not that the drug is unapproved in India. Note what was deliberately not done: the absence was established from Indian sources only, and no American or European value was consulted to produce any part of it. A jurisdiction's emptiness is established in that jurisdiction, never inferred from another's fullness.

PITFALL

Counting search hits instead of reading them. The regulator-domain search returned five documents and none names esketamine; a count would have recorded five findings where the correct number is zero. The error runs the other way in the trial registry, where a zero was believable only because a **positive control** had returned a known record first. Open what the search returns, and prove the instrument can see before reporting that it saw nothing.

2.7 Reading a regulatory record without being fooled by it

Three failure modes turned up while this position was assembled, each producing a confident, well-cited, wrong sentence.

- **Corroboration that is circular.** The European indication was first taken from the agency's published medicines register and checked against the public assessment report page. The register is a machine-generated dump of those same pages, so the corroboration was circular and blind to precisely the failure that had occurred, a stale summary layer. Annex I section 4.1 of the Summary of Product Characteristics is the genuinely independent instrument and is the authority for the indication text; the register and the report page corroborate product identity and authorisation status only.
- **Truncation that closes with a full stop.** A summary reproducing the first part of an indication and stopping at a sentence boundary reads as complete, because nothing signals the loss. The European row was tested for it from the start; the American row was not, and the stored copy of this very label stopped mid-clause at "does not preclude the need ", hiding a whole third limitation until an author refused to complete it from memory. The instrument was then pointed at the American column too, and the repaired row passed at 1649 characters against 1649. Quote an indication to its end, and where the copy in hand is cut off, say so and quote only the part that is whole.
- **A date wearing the wrong field name.** The register's decision-date column carries the latest European Commission decision, not the first; the build proved it on another product in the same workbook, where the column showed a 2026 date against an authorisation granted in 2014. Label a date for what it is: initial marketing authorisation, latest decision, earliest approval entry in an application record. A true date of one thing published under the field name of another is factually sourced and substantively false.

All three share one shape: the instrument returned something, the reader accepted it, and nobody asked what it was incapable of showing. The discipline is not what did the source say, but what could it not have told me.

2.8 What a complete answer contains

Reduced to four sentences:

- United States, on the label effective 13 March 2026, read 3 August 2026: treatment-resistant depression in adults, as monotherapy or in conjunction with an oral antidepressant; and depressive symptoms in major depressive disorder with acute suicidal ideation or behavior, in

conjunction with an oral antidepressant, under the three Limitations of Use quoted whole above.

- European Union, on the Summary of Product Characteristics read 1 August 2026 and re-read verbatim the next day: two authorised indications, treatment-resistant major depressive disorder in combination with an SSRI or SNRI, and the acute short-term reduction of depressive symptoms constituting a psychiatric emergency, co-administered with oral antidepressant therapy. Neither is a monotherapy indication.
- India, as of 1 August 2026: no CDSCO marketing approval record located, across 66 published approval lists, a verified regulator-domain search and a positive-controlled trial-registry search. A declared gap, not a statement that the drug is unapproved.
- The monotherapy question therefore has three answers and no single one.

MNEMONIC

Where, What, Who says so, When

The four fields of a status claim, in the order an examiner checks them. Where is the jurisdiction; What is the indication in the source's own words, whole and to its end; Who says so is the named regulator and the named document, never "approved internationally"; When is the date the status was read. An answer missing the fourth is an undated assertion in a domain that changes.

Esketamine permits monotherapy in one jurisdiction, forbids it in another, and has no located record in a third: the monotherapy question has no answer until a regulator and a date are named.

3 · Zuranolone and neurosteroid antidepressants for postpartum and major depression

This is the class that turned an endogenous hormone into a fourteen-day prescription, and then discovered how narrow the licence for it would be. Neuroactive steroids are the first antidepressants reasoned backwards from obstetric physiology rather than forwards from a monoamine, and two reached a regulator: brexanolone, an infusion now marketing status Discontinued in the United States, and zuranolone, the oral successor. The marks sit in three places: the subunit, the clock and the border.

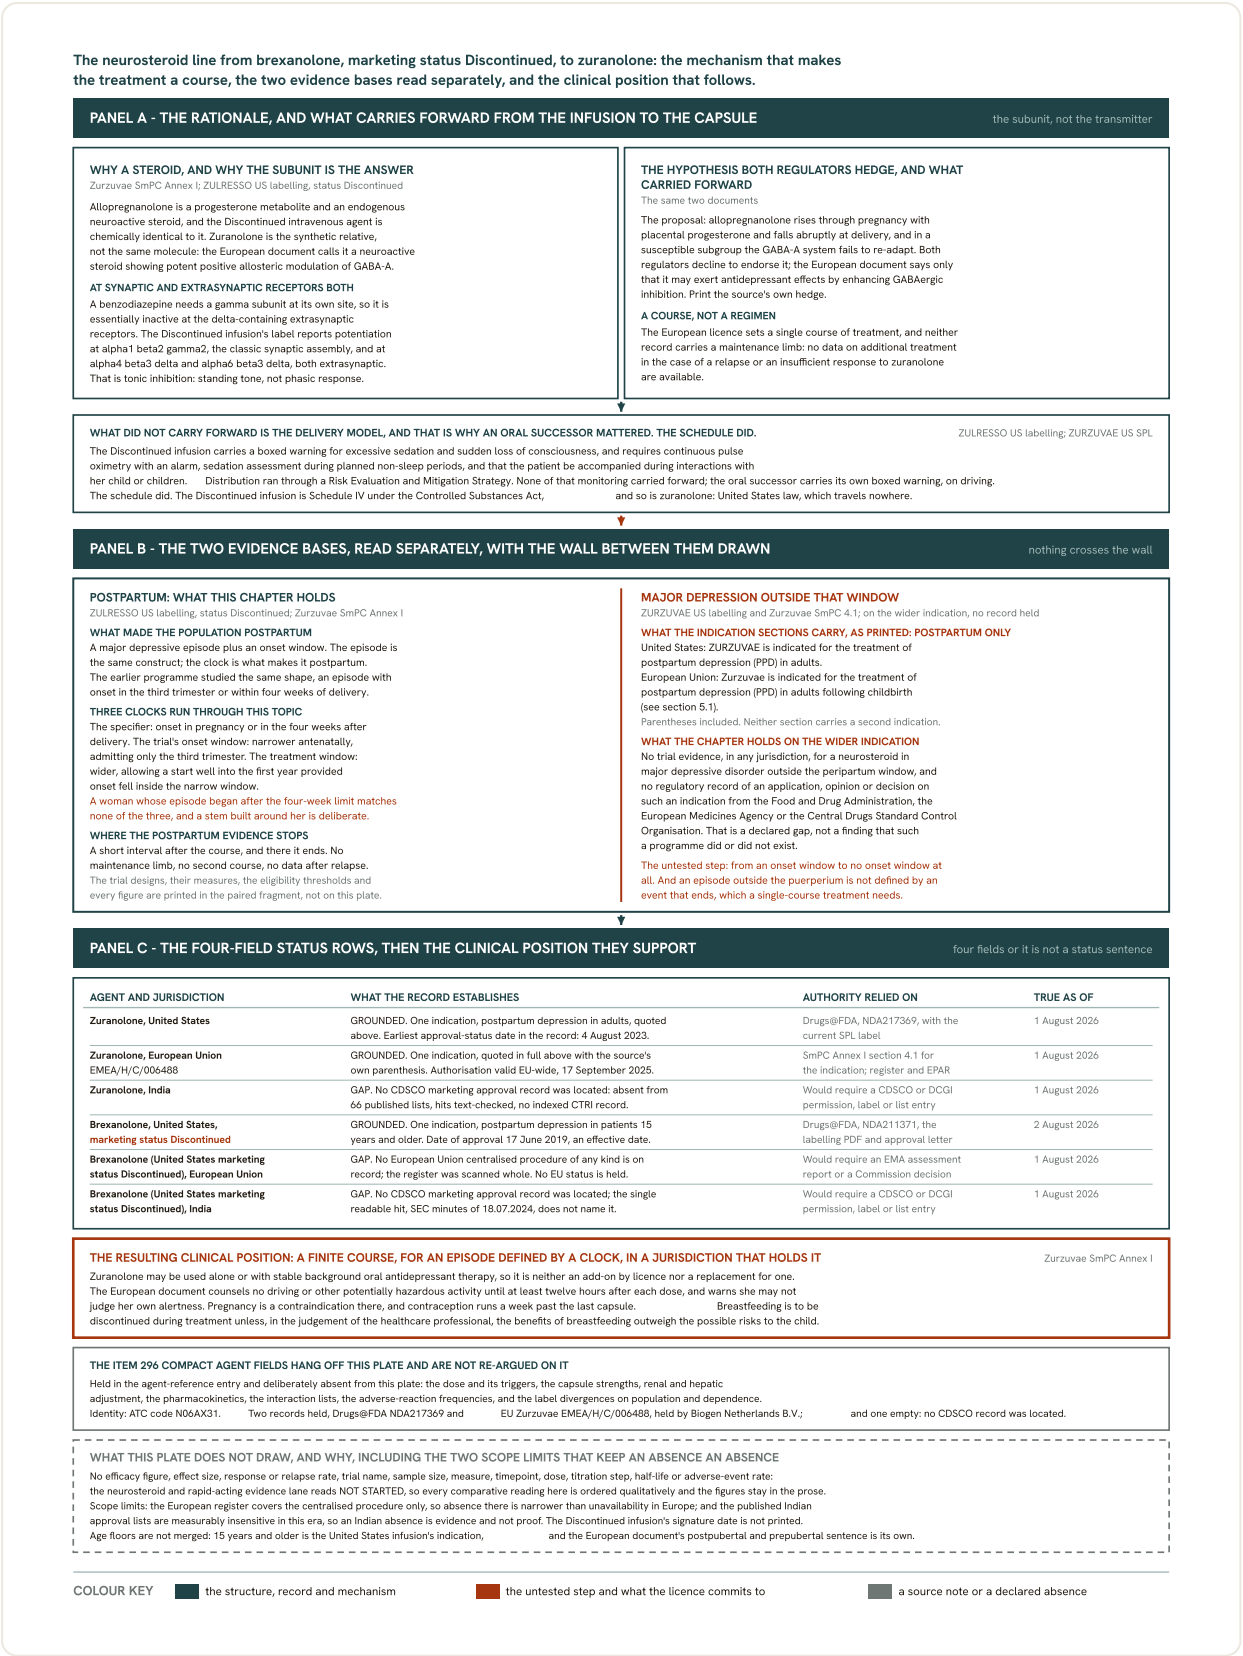


Fig 38.2 Zuranolone and the neurosteroid line from brexanolone, marketing status Discontinued: the mechanism, the two evidence bases read separately, and the position that follows. What carried forward from the Discontinued infusion is a single course of treatment with no maintenance limb, and the United States schedule; what did not is the monitoring service it required. Each indications section names one indication: postpartum depression (PPD) in adults in the United States, and postpartum depression (PPD) in adults following childbirth (see section 5.1) in the European Union. (Zurzuvae EMA EPAR Annex I SmPC; ZURZUVAE and ZULRESSO United States labelling, the latter marketing status Discontinued; the EU medicines register; the published CDSCO approval lists.)

3.1 Why a steroid, and why after delivery

Allopregnanolone is a progesterone metabolite and an endogenous neuroactive steroid modulating inhibitory transmission. The United States label for the intravenous agent: brexanolone is a neuroactive steroid GABA-A receptor positive modulator chemically identical to endogenous allopregnanolone, marketing status Discontinued. Zuranolone is the synthetic relative, not the identical molecule. Its European Summary of Product Characteristics calls it a **neuroactive steroid** showing potent **positive allosteric modulation** of the GABA-A receptor, enhancing GABAergic activity at **synaptic and extrasynaptic GABA-A receptors** and increasing their cell surface expression in vitro.

That last clause is the examinable difference. The benzodiazepine site needs a gamma subunit, so benzodiazepines act at synaptic receptors and are essentially inactive at the delta-containing extrasynaptic receptors carrying the persistent background current. Neuroactive steroids bind elsewhere and modulate both. The United States brexanolone label, marketing status Discontinued, shows it: the drug potentiated GABA-mediated currents from recombinant human receptors expressing alpha1 beta2 gamma2 subunits, the classic synaptic assembly, and alpha4 beta3 delta and alpha6 beta3 delta subunits, which are extrasynaptic. Modulating the extrasynaptic population modulates **tonic inhibition**, the standing tone rather than the phasic response. An answer saying only "it works on GABA" has given away the subunit, the part being tested.

The physiology that pointed the class at the puerperium is a hypothesis. Allopregnanolone rises through pregnancy with placental progesterone and falls abruptly at delivery; the proposal is that in a susceptible subgroup the GABA-A system fails to re-adapt. Both regulators decline to endorse it. The Discontinued intravenous agent's label says the mechanism of brexanolone in postpartum depression is not fully understood but is thought to relate to positive allosteric modulation of GABA-A receptors, and the European zuranolone document says only that it may exert antidepressant effects by enhancing GABAergic inhibition. Print the source's own hedge.

RULE **The mechanism is why this is a course and not a regimen.** The United States licence for brexanolone, Drugs@FDA NDA211371, approved for postpartum depression on 17 June 2019 and now marketing status Discontinued, administers one 60 hour infusion. The European licence for zuranolone, Zurzuvae, EMEA/H/C/006488, authorised for postpartum depression on 17 September 2025, sets a single 14 day course. Neither document carries a maintenance schedule, which is why every question about continuation, relapse and a second course meets an absence of evidence, not a dosing table.

3.2 Brexanolone, the Discontinued first-in-class, and its delivery model

Brexanolone, marketing status Discontinued, proved the class could work. Its United States labelling indicates it for postpartum depression in patients 15 years and older. Administration is a **continuous intravenous infusion over 60 hours**: 30 micrograms per kilogram per hour for the first 4 hours, 60 for the next 20 hours, 90 for the following 28 hours, then 60 for 4 hours and 30 for the last 4 hours, holding at 60 through the middle block if 90 is not tolerated. The shape: open low, hold high, taper, never stopped abruptly.

The monitoring turned it into a service. This Discontinued product's label carries a boxed warning for excessive sedation and sudden loss of consciousness, and requires continuous pulse oximetry with an alarm, sedation assessment every 2 hours during planned non-sleep periods, and that patients be accompanied during interactions with their child or children while the infusion runs. Distribution ran through a **Risk Evaluation and Mitigation Strategy**, and the product is Schedule IV under the Controlled Substances Act. In the trials behind this Discontinued agent's licence, sedation or somnolence forced dose interruption or reduction in 5 per cent of treated patients against none on placebo, loss or altered state of consciousness occurred in 4 per cent against none on placebo, and recovery took 15 to 60 minutes. Most services cannot deliver that whatever the effect size, which is the case for an oral successor.

TRAP **A true date wearing the wrong field name.** The United States date of approval for brexanolone, marketing status Discontinued, is **17 June 2019**, while the regulator's submission-status field for the original application holds 19 March 2019. Both are real, and the approval letter explains the gap: because controlled-substance scheduling was incomplete, it states the date of approval shall be the date the Drug Enforcement Administration publishes a Federal Register notice announcing interim final scheduling. A date is not a fact until you know which event it timestamps.

3.3 Zuranolone: the shape of a fourteen-day treatment

The European posology sets **50 mg once daily for 14 days** as a **single course of treatment**, reducible to 40 mg once daily for 14 days if not tolerated, and states dosing may be stopped without down-titration. There is no maintenance limb and no second course: the same document records that no data on additional treatment in the case of a relapse or an insufficient response to zuranolone are available.

The capsules are taken in the evening with fat-containing food, as either a meal or a snack, an exposure requirement, since absorption is food-dependent. A missed dose is not made up. Zuranolone may be used alone or with stable background oral antidepressant therapy, unlike agents licensed only as add-ons.

Zuranolone is metabolised primarily by CYP3A. Strong inhibition raises exposure, so the dose drops to 30 mg once daily through the 14 day treatment period when a strong CYP3A inhibitor is co-prescribed, and grapefruit is avoided for the same reason. Induction does the opposite: exposure is reduced by 85 per cent in the presence of rifampicin, which is why inducers are avoided outright rather than dose-adjusted. In an Indian setting that list is not hypothetical: carbamazepine, phenytoin, phenobarbital and rifampicin. Renal and hepatic adjustment and the full adverse-reaction table sit in the agent-reference entry.

3.4 The postpartum evidence, read the way an examiner reads it

Two development programmes, one population, and no licence to mix their numbers.

PROGRAMME	WHO WAS STUDIED	PRIMARY ENDPOINT	WHAT IT SHOWED
Brexanolone, status Discontinued: two placebo-controlled studies, Studies 1 and 2	Women aged 18 to 45 meeting DSM-IV criteria for a major depressive episode, onset in the third trimester or within 4 weeks of delivery; Study 1 severe, HAM-D total at or above 26, Study 2 moderate, HAM-D total between 20 and 25	Change from baseline in HAM-D total at the end of infusion, the sixtieth hour; the thirtieth day pre-specified as secondary	Severe illness: the 90 microgram per kilogram per hour target beat placebo by 3.7 points and the 60 microgram target by 5.5; moderate illness: the 90 microgram target by 2.5
Zuranolone: one randomised, double-blind, placebo-controlled multicentre study, 217-PPD-301, NCT04442503, SKYLARK	Women aged 18 to 45 meeting DSM-5 criteria for a major depressive episode, baseline HAMD-17 at or above 26, onset limited to the third trimester or within 4 weeks of delivery, treatment started up to 12 months after childbirth, followed 4 weeks after the 14 day course	Change from baseline in HAMD-17 at the fifteenth day, the day after the course ends	Least-squares mean change minus 15.6 against minus 11.6, a treatment difference, 4.0 points, 95 per cent confidence interval 1.7 to 6.3, p equals 0.0007

Speed: separation is present by the third day, at 3.4 points, the property separating this class from a conventional antidepressant. Durability inside the trial: the difference is still measurable at the forty-fifth day, at 3.5 points, a month after the last capsule, where follow-up stops. The placebo arm improved by nearly 12 points, so the drug effect is a modest increment on a large non-specific recovery, as a naturally remitting condition with intensive trial contact produces.

4.0	57.0%	38.9%	14
TREATMENT DIFFERENCE IN HAMD-17 POINTS	RESPONDERS ON ZURANOLONE	RESPONDERS ON PLACEBO	DAYS, THE WHOLE COURSE

Response, at least a halving of the baseline total, was reached by 57.0 per cent against 38.9 per cent on placebo, a relative risk 1.409 and an absolute difference, 16.4 percentage points. Among those responders, three subjects, 5.7 per cent, relapsed by the trial's definition of two consecutive totals at or above 20 after the fifteenth day, and none met the definition of rebound. Two cautions: the denominator is small, so the confidence around 5.7 per cent is wide, and the window closes four weeks after the course.

3.5 What happened to the major depressive disorder claim

Most answers go wrong here by writing a sentence that is nearly true. Both regulatory records held for zuranolone name one indication, and it is not major depressive disorder. The United States labelling reads that ZURZUVAE is indicated for the treatment of postpartum depression (PPD) in adults, carried twice in that record. The European therapeutic-indications section, Annex I section 4.1, reads that Zurzuvae is indicated for the treatment of postpartum depression

(PPD) in adults following childbirth (see section 5.1). Both are quoted as printed, parentheses included. Neither indications section carries a second indication.

Set that against the trial population. Every patient in the pivotal study met criteria for a **major depressive episode**; what made the study postpartum was the clock, onset restricted to the third trimester or the 4 weeks after delivery. The untested step is from that onset window to no onset window at all.

Three clocks run through this topic; merging them produces a wrong sentence. The DSM-5 peripartum-onset specifier covers onset during pregnancy or in the 4 weeks after delivery. The trial's onset window is narrower antenatally, admitting only the third trimester. The treatment window is wider, since subjects could start treatment up to 12 months after childbirth provided symptoms began inside the narrow window. A woman presenting at nine months with an episode that began at eight matches none of the three, and a stem built around her is deliberate.

TRAP **Reading the postpartum evidence as depression evidence.** This chapter holds no trial evidence, in any jurisdiction, for a neurosteroid in major depressive disorder outside the peripartum window, and no regulatory record of an application, opinion or decision on such an indication from the United States Food and Drug Administration, the European Medicines Agency or the Central Drugs Standard Control Organisation. That is a declared gap in this chapter's evidence, not a finding that such a programme did or did not exist. Writing "zuranolone is a new rapid-acting oral antidepressant for depression" crosses the indication boundary and the evidence boundary at once.

One further argument comes from the shape of the licence, not from a decision this chapter lacks: the whole authorised European exposure is a single fourteen-day course, with no data on treatment after relapse. Major depression outside the puerperium is not defined by an event that ends, so a treatment with no second course and no maintenance limb has a structural problem in that indication whatever its acute effect.

3.6 Sedation, the driving rule, and the abuse question

Every frequency here is from the European Summary of Product Characteristics. The most frequent adverse reactions were somnolence at 27.6 per cent, dizziness at 13.3 per cent and sedation at 11.2 per cent; the reaction classed as serious was confusional state, at 1.3 per cent. Only 2 per cent discontinued for an adverse reaction, while 14.3 per cent had a dose reduction or interruption. Among those on the higher dose, 69.7 per cent of somnolence events and 58.3 per cent of sedation events occurred within the first 2 days of treatment. Warn her that the first two evenings are the heaviest.

The same European document counsels patients not to drive or engage in other potentially hazardous activities until **at least 12 hours after taking each dose**, and advises they may not be able to judge their own ability; the United States labelling carries the same restriction in a boxed warning. Two driving-simulation studies found dose-dependent impairment nine hours after dosing, and modelling of the lateral-position data put median placebo-corrected impairment

below the threshold for a 0.05 per cent blood alcohol concentration only by the twelve-hour mark.

PITFALL

Treating the sedation as a side effect to be tolerated rather than as the thing structuring the prescription. The patient is by definition caring for an infant, and the two documents handle that differently. The United States label of the Discontinued intravenous product requires a companion during interactions with the children. The European zuranolone document names driving and hazardous activities and does not name infant care among them, so an answer must not attribute to zuranolone a chaperone requirement the European document does not name. The counselling point: she may not judge her own alertness reliably, and night-time infant care over the first evenings is a conversation to have before the first capsule.

Abuse potential is assessed in that same European document. In a human abuse potential study in 60 recreational users of central nervous system depressants, zuranolone at 30, 60 and 90 mg showed dose-dependent abuse potential against alprazolam at 1.5 mg and 3 mg on positive subjective measures including drug liking and good drug effects, while clinical-trial data indicate low physical dependence potential. Liking is not dependence. Caution is advised where there is a history of abuse of or addiction to alcohol or other substances.

3.7 Pregnancy, contraception and breastfeeding

In the European document zuranolone is contraindicated during pregnancy, on animal reproductive toxicity and an absence of human data, and women of childbearing potential must use effective contraception during treatment and for 7 days following discontinuation. Contraception is part of the prescription, and cover runs a week past the last capsule.

The numbers and the instruction pull against each other; print both. A lactation study measured the **relative infant dose** under 1 per cent: the mean was 0.357 per cent at a 30 mg maternal dose, and simulation put the expected mean at 0.738 per cent and 0.984 per cent for a 50 mg maternal dose at two assumed milk intakes. Despite figures that low, the European document instructs that breastfeeding should be discontinued during treatment unless, in the judgement of the healthcare professional, the benefits of breastfeeding outweigh the possible risks to the child, because the effect on a breastfed infant is unknown and milk-production data are limited. The United States label for the Discontinued intravenous agent records a higher relative infant dose, 1 to 2 per cent of the maternal weight-adjusted dosage, with oral bioavailability under 5 per cent keeping infant exposure low. Quantified low exposure and a precautionary instruction are not a contradiction.

The age boundaries differ and must not be merged. The European document records that for zuranolone, safety and efficacy in postpubertal females under 18 years old have not been established, and there is no relevant use in prepubertal females. For the Discontinued intravenous agent the United States indication runs to patients 15 years and older.

GOOD TO REMEMBER

Two questions decide relevance, and neither is pharmacological: when the episode began, and whether she is breastfeeding. The first decides whether any evidence held here applies to her; the second, whether the licensed instruction can be followed without imposing a loss she may not accept.

3.8 Status by jurisdiction: what each regulator's record actually says

Every regulatory sentence here carries four fields, and one missing a field is not a status sentence: the jurisdiction, the exact indication or purpose, the authority the claim rests on, and the date the status was true. The interesting rows are the empty ones.

AGENT AND JURISDICTION	WHAT THE RECORD SAYS	AUTHORITY RELIED ON	STATUS TRUE AS OF
Zuranolone, United States	Indicated for the treatment of postpartum depression (PPD) in adults. Earliest approval-status date in the record, 4 August 2023	Drugs@FDA, NDA217369	1 August 2026
Zuranolone, European Union	Indicated for the treatment of postpartum depression (PPD) in adults following childbirth (see section 5.1); authorised as Zurzuvae, product number EMEA/H/C/006488, holder Biogen Netherlands B.V.; **marketing authorisation** valid throughout the European Union from **17 September 2025**	Indication from Annex I section 4.1 of the Summary of Product Characteristics; identity and status from the medicines register and assessment-report page	1 August 2026
Zuranolone, India	No CDSCO marketing approval record was located. Absent from all 66 published approval lists; the three documents a CDSCO-domain search returned were text-checked and none names the drug; no indexed CTRI record	Would require a CDSCO or DCGI new-drug permission, an approved label or an official list entry	1 August 2026
Brexanolone, United States, marketing status Discontinued	Indicated for the treatment of postpartum depression (PPD) in patients 15 years and older [see Clinical Studies (14)]; date of approval 17 June 2019; marketing status Discontinued	Drugs@FDA, NDA211371; indication from the single-column prescribing information, date from the approval letter	2 August 2026
Brexanolone (United States marketing status Discontinued), European Union	No European Union centralised procedure of any kind is on record for this active substance; the register was scanned whole and returns no row under any name. No European marketing status is on record	Would require an EMA assessment report or a Commission decision	1 August 2026

AGENT AND JURISDICTION	WHAT THE RECORD SAYS	AUTHORITY RELIED ON	STATUS TRUE AS OF
Brexanolone (United States marketing status Discontinued), India	No CDSCO marketing approval record was located. Absent from all 66 published approval lists; the single CDSCO-domain hit, expert-committee minutes dated 18 July 2024, is readable and does not name the substance; no indexed CTRI record. No Indian marketing status is on record	Would require a CDSCO or DCGI new-drug permission, an approved label or an official list entry	1 August 2026

Two properties must be stated in words or a reader will over-read the table. The register covers the European Union centralised procedure only, so a medicine authorised nationally through the decentralised or mutual-recognition route would not appear in it. What the empty row establishes is that no central application was ever concluded, by authorisation, refusal, withdrawal, lapse, revocation or suspension, since the register carries all those outcomes. That is narrower than "not available anywhere in Europe". For this substance the scope limit bites hard, and the confirming second instrument never ran on any gap row.

India second. The published CDSCO approval-list corpus is demonstrably incomplete, and no positive control establishes its sensitivity in the period these agents belong to: the one drug in this chapter whose Indian approval is known from outside the corpus, through expert-committee minutes, appears nowhere in the 66 documents, and the list declaring full-year coverage for that year carries entries dated after the approval it omits. A corpus that misses a drug the same regulator says it approved cannot prove another drug was not approved. So the sentence a candidate may write is not that these agents are unapproved in India, but that **no CDSCO approval record was located** as of August 2026, naming what was searched, on the authority of the Central Drugs Standard Control Organisation and the Drugs Controller General of India.

IN INDIA

Four things are mistaken for an Indian approval record. A CTRI registration records that a trial was registered, not that a drug was approved. Expert-committee minutes are advisory proceedings. A brand in a pharmacy list or a distributor catalogue is a commercial artefact. And an approval in Washington or Amsterdam is a fact about Washington or Amsterdam. The record statement: as of August 2026 no CDSCO approval record was located for either agent. This chapter holds no instrument that measured Indian availability either way.

- Mechanism: a neuroactive steroid modulates synaptic and extrasynaptic GABA-A receptors, reaching a tonic inhibition a benzodiazepine does not.

- A course, not a regimen: fourteen days orally, sixty hours by infusion, no second course described for either.
- A major depressive episode with an obstetric onset window, which is why the licence reads postpartum depression and the evidence goes no further.
- A status sentence names a jurisdiction, an exact indication, an authority and a date, or it is not written.

MNEMONIC

Subunit, Course, Clock, Border

Subunit separates this class from the benzodiazepines. Course is the finite exposure that follows. Clock is the onset window making a major depressive episode a postpartum one. Border is the rule that a regulator's record speaks only for its own territory.

The mechanism explains why the treatment is short, the onset window explains why the licence is narrow, and the jurisdiction explains why a true sentence about this drug can still be the wrong answer.

4 · Rapid-acting antidepressants: mechanism class and clinical positioning

A rapid-acting antidepressant is defined by a claim about time and licensed by a sentence about an indication. The class lives between those two objects. Its members share a response measured in hours to days rather than weeks, produced outside monoamine reuptake, and nothing else: not the receptor, the setting, the population each licence names, how long the treatment runs, or whether any regulator has concluded anything. Marks go in three places. Candidates write **rapid-acting antidepressant** as though a regulator had licensed the category; they collapse five members and three jurisdictions into one verdict; and they read the silence of one search as a fact about five products. The clinical accounts belong elsewhere: ketamine and esketamine in treatment-resistant depression to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes, their mood-disorder positioning to the Mood disorders: pharmacotherapy notes.

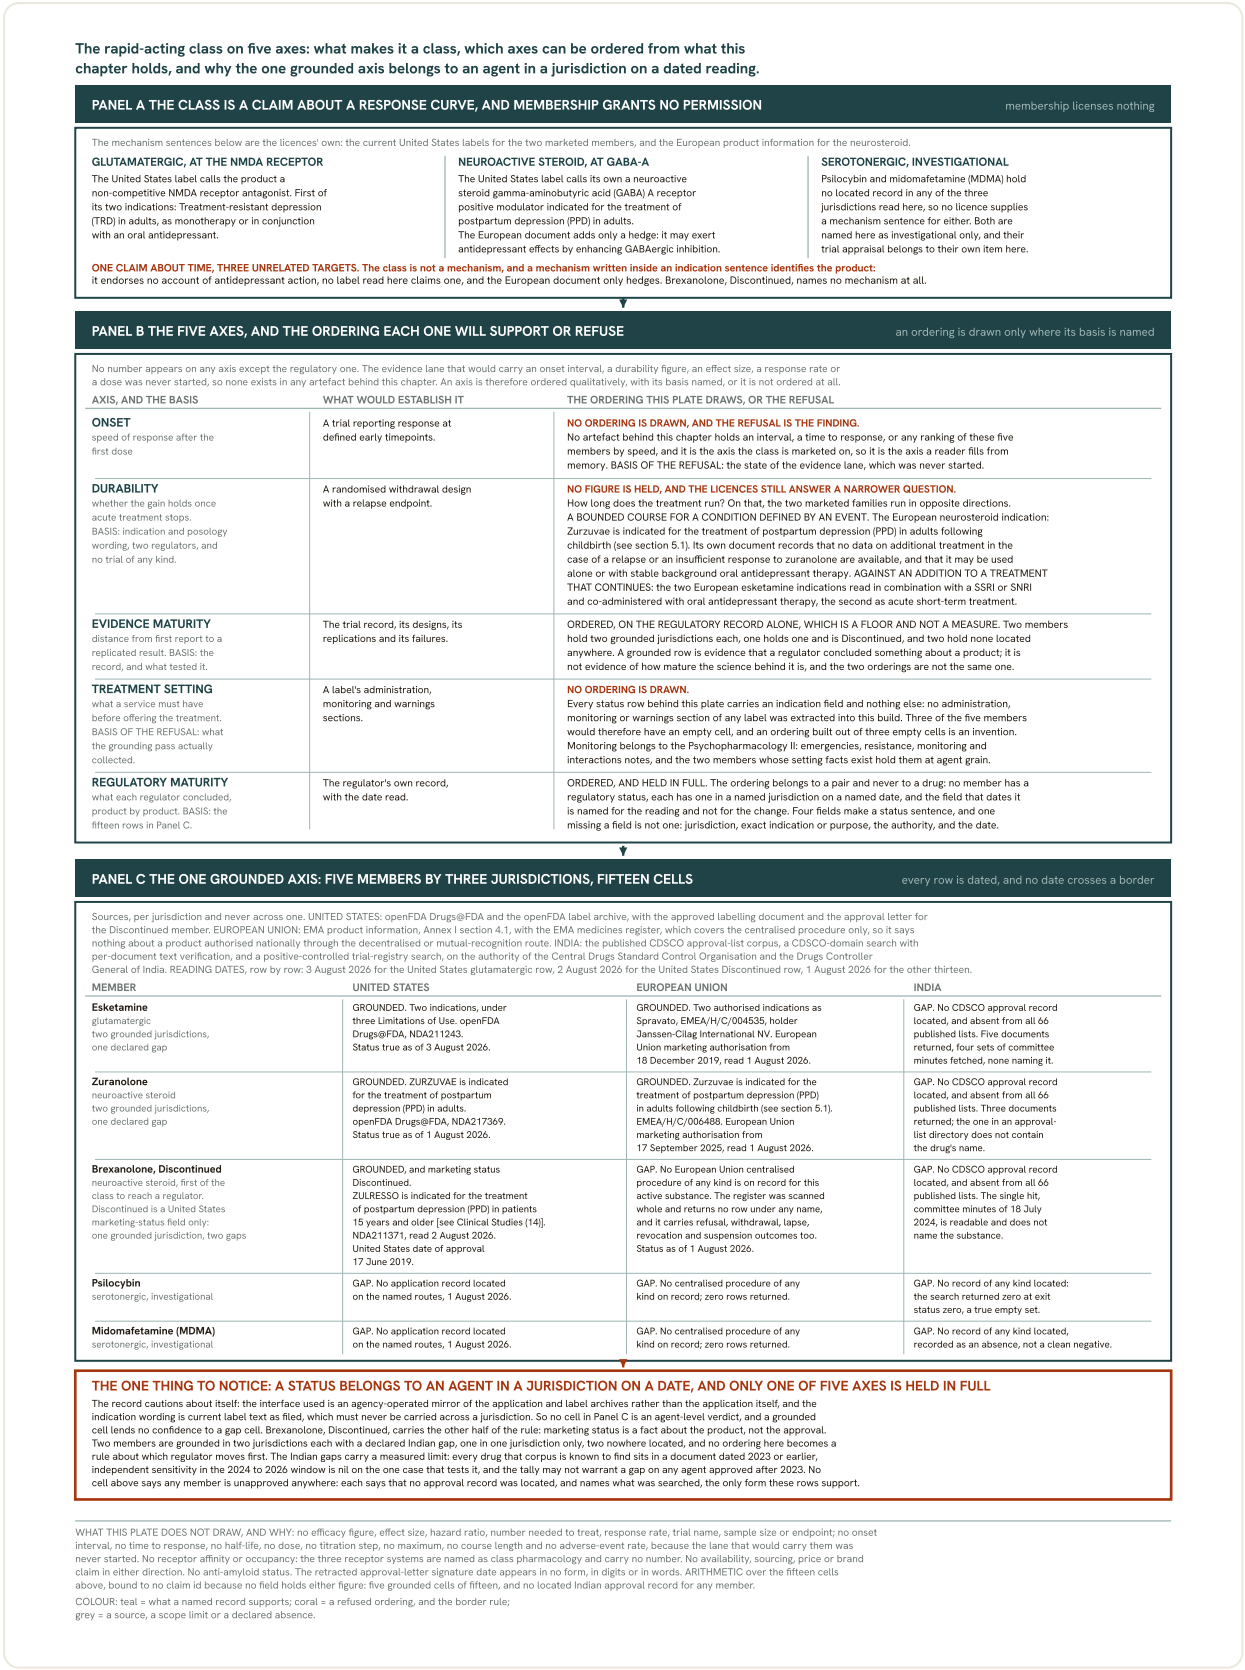


Fig 38.3 The rapid-acting class on five axes: only regulatory maturity is held in full, three of five not at all. Onset and treatment setting are refused orderings, each basis named. Durability is licence wording alone: a bounded course for a condition defined by an event, in adults following childbirth (see section 5.1), against an addition to continuing treatment, in combination with a SSRI or SNRI. Regulatory maturity belongs to an agent in a jurisdiction on a dated reading: five grounded cells of fifteen, with 17 June 2019 the United States date of approval for brexanolone, marketing status Discontinued. (openFDA Drugs@FDA, label archive and approval letter; EMA Annex I section 4.1, centralised procedure only; CDSCO approval lists and domain search, on CDSCO and DCGI authority.)

4.1 What makes this a class, and what the class promises

No regulator in these records authorises a rapid-acting antidepressant. Every authorisation names a product, an indication and a population, so the class is a construct built for a clinical purpose and carries no permissions. A sentence travelling from "this agent is rapid-acting" to where it may be used, in whom, or on whose authority has become a permission without passing through a document.

The promise sets the two questions the class is judged on. Conventional practice tolerates a latency of weeks and builds the service around that wait. Shorten it and two questions open: what happens to the gain when acute treatment ends, and what was the shortened wait supposed to buy?

One regulator has answered the second in writing. Among the **Limitations of Use** on the current United States label for esketamine stands the statement that The effectiveness of SPRAVATO in preventing suicide or in reducing suicidal ideation or behavior has not been demonstrated. The same label indicates the product for Depressive symptoms in adults with major depressive disorder (MDD) with acute suicidal ideation or behavior in conjunction with an oral antidepressant. The licence covers a population defined by suicidal ideation while declining any effect on the event that population is at risk of. Symptom change and prevention of an outcome are separate endpoints, and the regulator separated them. An answer arguing from speed of response to a suicide-prevention benefit makes the leap the licence refuses.

Notice what kind of negative that is, because a candidate trained to distrust categorical negatives may mistake it for one. It states what a trial programme demonstrated, printed inside a licence by the authority that granted it, so it is a positive fact about an evidence base; a claim that a drug is unapproved somewhere needs a regulator's affirmative statement and appears nowhere here without one. The same label's SPRAVATO is not approved as an anesthetic agent is again a licence describing its own boundary.

RULE **Rapid-acting describes a response curve and never grants a permission.** Every sentence about what may be given, to whom, where and on whose authority must be sourced to one named regulator's record for one named product, with the date read. The class name cannot stand in for a document.

4.2 Three mechanistic families, and the mechanism sentence a licence gives you

Three receptor systems sit under one heading: the glutamatergic agents at the NMDA receptor complex, the neuroactive steroids as positive modulators at the GABA-A receptor, and the serotonergic psychedelics, investigational throughout, at serotonin 2A receptors. One claim about time, three unrelated targets: this class is not a mechanism, and an answer offering one has answered another question.

Two of the three have a mechanism written into the licence. Esketamine's United States label calls the product **a non-competitive NMDA receptor antagonist**, spelling the receptor name out in full, indicated for Treatment-resistant depression (TRD) in adults, as monotherapy or in conjunction with an oral antidepressant. Zuranolone's United States label calls its own **a**

neuroactive steroid gamma-aminobutyric acid (GABA) A receptor positive modulator indicated for the treatment of postpartum depression (PPD) in adults. A mechanism inside an indication sentence identifies the product; it endorses no account of antidepressant action, and no label read here claims one.

The third teaches it from the other side. Brexanolone, whose United States **marketing status** is Discontinued, carries an indication naming no mechanism: ZULRESSO is indicated for the treatment of postpartum depression (PPD) in patients 15 years and older. The neuroactive-steroid account of it comes from the literature, not the licence. Asked what a licence says, answer only that.

No receptor affinity, occupancy figure, onset interval, half-life or comparison of potency or speed between these families is held by any artefact behind this chapter, because the evidence lane that would carry them was never completed. Writing one from memory would be a fabricated citation, the one defect no later pass repairs.

4.3 The five axes, and which of them anything here fills

A class map is useful only if it is honest about its columns, and here one of five axes is held in full, one narrowly, and three not at all. A table hiding that invites a reader to fill the empty columns from memory and quote the result as though the map supplied it.

AXIS	WHAT IT WOULD MEASURE	WHAT WOULD ESTABLISH IT	WHAT THESE NOTES HOLD
Onset	Speed of response after the first dose	A trial reporting response at defined early timepoints	Nothing: no interval, no time to response, no ranking by speed
Durability	Whether the gain holds once acute treatment stops	A randomised withdrawal design with a relapse endpoint	No figure. Only licence wording, separating a bounded course from a continuing treatment
Evidence maturity	Distance from first report to replicated result	The trial record, its designs, replications and failures	Only the regulatory record and the instruments that tested it: a floor, not a measure
Treatment setting	What a service must have before offering the treatment	A label's administration, monitoring and warnings sections	Nothing directly; never extracted. Monitoring belongs to the chapters named above
Regulatory maturity	What each regulator concluded, product by product	The regulator's own record, with the date read	Held in full, at the grain of one agent in one jurisdiction

The columns a candidate most wants to fill are the two the class is marketed on, so the defensible answer runs the map backwards: from what each regulator concluded, to what the licence implies about treatment length, then naming the rest unestablished.

4.4 Durability: the axis two licences answer in opposite ways

Nothing here quantifies durability, and the licences still answer the question. They answer the narrower version a service plans around: not how long a response lasts, but how long the treatment runs. On that the two marketed families run in opposite directions in the indication text.

The neurosteroid licences are built around a bounded situation. The European authorisation reads Zurzuvae is indicated for the treatment of postpartum depression (PPD) in adults following childbirth (see section 5.1). The United States indication for the same molecule reads ZURZUVAE is indicated for the treatment of postpartum depression (PPD) in adults. A condition defined by its relation to an event implies a treatment with an end, and the European text names the event where the American does not.

The glutamatergic member's European licence runs the other way, indicating esketamine in combination with a SSRI or SNRI for treatment-resistant major depressive disorder, and in its second authorised indication co-administered with oral antidepressant therapy for a moderate to severe episode as acute short-term treatment, for the rapid reduction of depressive symptoms, which according to clinical judgement constitute a psychiatric emergency. Two facts sit there: both European indications license the agent inside a continuing antidepressant treatment rather than instead of one, and the second carries its own durability limit in the words acute short-term treatment.

The class therefore splits along a line unrelated to receptor: a bounded course for a bounded condition, against an addition to a treatment that continues. Both say how long a prescription runs, not how long a response lasts, which decides staffing and review intervals. The United States monotherapy line belongs to esketamine's own section here.

One demographic comparison is available. Brexanolone, marketing status Discontinued, is licensed, in the United States, in patients 15 years and older, while the indications quoted above for both marketed members name adults. The class's only reach below adulthood sits on the member that is Discontinued, a fact about a licence rather than an option to offer anyone.

> **GOOD TO REMEMBER**
>
> Two questions answer the planning problem from the licence alone, without a trial: is this agent licensed as an addition to something that continues, and is the condition it names defined by an event that ends? Different clinical commitments, and the patient should be told which is starting.

4.5 Regulatory maturity belongs to a pair, not a drug

Four fields make a status sentence and one missing a field is not one: the jurisdiction, the exact indication or purpose in the source's words, the authority the claim rests on, and the date it was true. That last field's name is worth using, **status true as of**, because it dates the reading and not the change. The United States rows here were read on 3 August 2026 for esketamine, 1 August for zuranolone and 2 August for brexanolone, marketing status Discontinued, and every European and Indian row on 1 August 2026.

Five members and three jurisdictions make fifteen cells, and they disagree inside one agent: no member has a regulatory status, each has one in a named jurisdiction on a named date. The record says so of itself: the interface used is an agency-operated mirror of the application and label archives rather than the application itself, and the indication wording is current label text as filed, which must never be carried across a jurisdiction.

MEMBER	GROUNDED RECORDS	DECLARED GAPS	WHAT MAY NOT BE WRITTEN
Esketamine, glutamatergic	United States, two indications under three Limitations of Use, NDA211243, read 3 August 2026. European Union, two indications as Spravato, EMEA/H/C/004535, holder Janssen-Cilag International NV, from **18 December 2019**	India: no CDSCO approval record located	Either indication carried across, and any Indian status either way
Zuranolone, neurosteroid	United States, postpartum depression in adults, NDA217369. European Union, the same with following childbirth added, as Zurzuvae, EMEA/H/C/006488, holder Biogen Netherlands B.V., from **17 September 2025**	India: no CDSCO approval record located	The European wording quoted as the American indication, and the register's decision column read as an authorisation date
Brexanolone, neurosteroid, marketing status Discontinued	United States only, postpartum depression in patients 15 years and older, NDA211371, date of approval **17 June 2019**, read 2 August 2026	European Union: no centralised procedure on record; India: no CDSCO record located	That its approval has lapsed, and any inference from Discontinued about the gap rows
Psilocybin, investigational	None located	United States, European Union and India, all three declared	Any categorical negative, and any upgrade of investigational status to an authorisation
Midomafetamine (MDMA), investigational	None located	United States, European Union and India, all three declared	The same two, plus transfer from the other investigational member

The investigational pair belongs to the psychedelic-assisted therapy item in this chapter, which owns their trial appraisal. They appear only because a class comparison that drops its immature members flatters the class.

TRAP **Collapsing the map to one verdict per agent.** Two members hold two grounded jurisdictions each with a declared Indian gap, one holds one grounded jurisdiction and two gaps, two hold none anywhere. "Esketamine is approved" is true of two records, unsupported for a third, and unmarkable because it names no regulator. Never write an approval verb without the jurisdiction, source and date beside it, and never let a grounded row lend confidence to a gap row.

4.6 Brexanolone, Discontinued: a probe miss, a date set by scheduling, and an empty key

One row has been wrong in every interesting way and repaired each time. It also carries the rule for every sentence about it: teach this agent and carry its marketing status in the same breath, since marketing status is a fact about the product and not about the approval, and this product is Discontinued while its approval stands.

The row read GAP, on a verdict of no record found, until multi-route probing retrieved it. The verdict was a probe miss and not an absence: the record carries no agency-standardised name block, so queries through generic name, substance name and brand name returned not-found while the query reading the product's active-ingredient field returned it. The record's own sentence transfers: a decided action is not a licence to ignore a record that turns up.

Then the date, a controlled-substance lesson in clerical dress. The United States date of approval for brexanolone, marketing status Discontinued, is 17 June 2019, and the approval letter did not set it by its own signing. The letter states that the date of approval shall be the date on which the Drug Enforcement Administration publishes a notice in the Federal Register announcing the interim final scheduling of brexanolone, marketing status Discontinued. The original submission carries an earlier status date, that signature and not the approval; no number for it is printed here.

The third repair is worth the most and looks like housekeeping. This row's canonical earliest-approval key is deliberately absent, because writing the submission's signature date under a field name meaning first approval would put a retracted number back into the field every consumer reads. Rarer than getting that right, it writes down the cost. The carry check reads only the European, Indian and canonical earliest-approval keys, so it cannot see where this row's approval date now lives; the row sits outside that check until it learns the key, and the pass's claim to compare every date field is no longer true of this file. A correct repair left a control blind, and the record says so.

It is also the one indication here no instrument could test. All three text instruments returned unresolved extraction rather than a pass, for a documented reason: with no name block the label archive cannot be queried by application number, so the indication came from the approved labelling document instead; the pass's own discipline is that a failed extraction is unresolved extraction and never a pass. How it was read matters too. The indication was taken from the single-column prescribing information, not the two-column highlights block, whose extract interleaves the columns. Interleaved columns produce fluent sentences no document contains.

5 of 15	0 of 5	3 of 3	2 of 2
GROUNDED CELLS, FIVE MEMBERS BY THREE JURISDICTIONS	INDIAN APPROVAL RECORDS LOCATED FOR THESE MEMBERS	BREXANOLONE, DISCONTINUED, US TEXT CELLS UNTESTED, NOT PASSED	EU ESKETAMINE INDICATIONS CARRIED AGAINST SOURCE

The fourth figure counterweights the third: The completeness instrument found two indications in the European source document and two on the row, and the instrument set was validated by planting the real historical defect on that row, a dropped second indication, and confirming the guard fired. A pass from an instrument tested against the failure it exists to catch is worth something.

TRAP **Reading an untested cell as a clean one.** Three cells here were never tested, because the source text could not be retrieved by the route the instrument uses, and are recorded as unresolved rather than as passes. A reader who scans a matrix for the word fail, finds none and calls the column sound has learned nothing. The mirror error costs as much: an unresolved text check does not undermine a status resting on the labelling document and the approval letter.

4.7 The Indian column, and what nine search hits established

Every member has the same Indian entry: no located CDSCO marketing approval record. At class level that is one instrument's result five times over, so all five cells carry that one search's sensitivity, which the record measures rather than assumes.

What was opened is the lesson. For esketamine, a domain search returned five documents, four sets of subject expert committee minutes for neurology and psychiatry and one international cell file; all four were fetched, all carry an extractable text layer, and none names the drug, so all were search-index false positives and not evidence. For zuranolone, three came back, and the one in an approval-list directory was downloaded and its text does not contain the drug's name, as with both sets of minutes. For brexanolone, marketing status Discontinued, the single hit, committee minutes dated 18 July 2024, was fetched, is readable, and does not name the substance. Nine hits, eight documents opened, nine false positives, the fifth esketamine document scored one without being fetched: a search hit is not a document naming your drug.

The two investigational members returned nothing, and even that is graded. The domain search returned zero results at exit status zero, a genuine empty result set rather than a failed query, and the finding is still recorded as the absence of any located record rather than a clean negative, which would overstate what an instrument with a measured miss rate and three failed positive controls delivers.

Two sensitivity facts belong in any Indian sentence about a recent agent, and the first has a name, **era mismatch**. Every drug this corpus is known to find sits in a document dated 2023 or earlier, and independent sensitivity in the 2024 to 2026 window is nil on the one case testing it; the positive-control tally may not warrant a gap on any agent approved after 2023.

The authority is named, and its substitutes with it. The claim would rest on the Central Drugs Standard Control Organisation and the Drugs Controller General of India, and committee minutes, trial no-objection certificates, international cell confirmations, CTRI registration, brand

listings, pharmacy availability and manufacturer copy are none of them marketing authorisation. A CTRI entry records that a trial exists, not permission to market. The rule closing the column sits in each of these rows: no India row here asserts that a drug is not approved in India; each asserts that no CDSCO approval record was located, and names what was searched.

IN INDIA

Write the Indian column in this order. No CDSCO marketing approval record was located for any member as of 1 August 2026, on the three instruments named above, and the corpus has never been shown to find an approval from 2024 onward. Do not write that any is not approved in India, which is stronger than the evidence supports, and do not write that any is approved, because nothing retrieved says so. For a family: there is no located approval record to point to, and licensing is a question for the regulator's record, not a brand listing.

4.8 Choosing across the class, and what a complete answer contains

Five statements answer it in whatever form it is asked.

- The class is a claim about a response curve, not a licensed category, so membership licenses nothing.
- Three unrelated receptor systems share one heading, and a licence naming a mechanism identifies a product rather than endorsing an account of action.
- Of five axes a class map needs, only regulatory maturity is held in full, durability only in the narrow form licence wording settles.
- Regulatory maturity belongs to an agent and a jurisdiction on a stated date; a grounded row never lends confidence to a gap row.
- The Indian column is five copies of one search result with a measured sensitivity: a statement about a search, never about a product.

MNEMONIC

Curve, Country, Course, Corroboration

Curve is what the class name describes. Country is the jurisdiction every status sentence must name and never leave. Course is durability in the form a licence answers, how long the treatment runs. Corroboration is the last question before quoting a cell: was it tested, tested against the failure it exists to catch, or merely unmarked?

The class name describes a curve and grants nothing; status belongs to an agent and a jurisdiction on a dated reading; and an untested cell is not a clean one.

5 · Dual orexin receptor antagonists (DORAs) and sodium oxybate in sleep psychiatry

Four agents that one sleep service may use in the same week hold, between them, no single shared pattern of national approval. Three **dual orexin receptor antagonists** and one oxybate show that regulatory status belongs to an agent inside a jurisdiction and never to a class. The

pharmacology is not rebuilt here: the orexin mechanism, the agents, dosing, adverse effects and insomnia prescribing belong to the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes, and narcolepsy, cataplexy, orexin biology and the disease-specific use of sodium oxybate belong to the Eating and sleep-wake disorders notes. What is dated and new is the regulatory record. Three errors cost most: one class verdict for four agents, an approval carried across a border, and a sentence read as complete because it ends in a full stop.

5.1 Four fields, and why this group needs them most

A status sentence is checkable only if it carries the jurisdiction, the indication in the regulator's own words, the named source, and the date the reading was true. The rule that follows governs the chapter: a row may be written only in the jurisdiction it was grounded in, and none crosses a boundary. A declared absence is written as an absence, naming what was searched and whose authority a claim would rest on; an inferred figure is not written at all.

At the grain of one agent in one country these four generate twelve records that disagree. Seven carry an affirmative located record and five are declared absences, and which is which follows neither the agent, the class, the year nor the indication; the tally is computed across these rows. The three orexin antagonists are licensed, where they are, for insomnia, and the oxybate for narcolepsy with cataplexy or the excessive daytime sleepiness of narcolepsy.

RULE **An insomnia approval is never transferred to narcolepsy or cataplexy, and a narcolepsy approval is never transferred to insomnia.** Both families act on the same sleep-wake machinery and are prescribed by the same clinician, which is why the licences get swapped in an answer. Quote the indication belonging to that product in that country.

5.2 The United States: four records, four dates, one reading date that is not the others

All four carry an affirmative American record, read from the openFDA Drugs@FDA interface with a **structured product label** named by identifier and effective date, current for three rows and not for the oxybate. Three resolve to one application; the fourth does not, which is 5.3.

AGENT	APPLICATION AND LABEL	INDICATION AS FILED, QUOTED	EARLIEST APPROVAL-STATUS DATE IN THE RECORD
Suvorexant	NDA204569, label 21306af5 effective 10 March 2025	BELSOMRA (suvorexant) is indicated for the treatment of insomnia characterized by difficulties with sleep onset and/or sleep maintenance.	13 August 2014
Lemborexant	NDA212028, label c82838dc effective 24 February 2025	DAYVIGO is indicated for the treatment of adult patients with insomnia, characterized by difficulties with sleep onset and/or sleep maintenance	20 December 2019
Daridorexant	NDA214985, label 10e24e38 effective 26 February 2026	QUVIVIQ is indicated for the treatment of adult patients with insomnia, characterized by difficulties with sleep onset and/or sleep maintenance	7 January 2022
Sodium oxybate	Four applications, two original; text re-read from label 3839d447 effective 1 July 2025, whose own application number is NDA021196, not the row's current label	Xyrem is indicated for the treatment of cataplexy or excessive daytime sleepiness (EDS) in patients 7 years of age and older with narcolepsy.	**17 July 2002**

Three differences are examinable. The three insomnia texts differ: two name adult patients and the earliest names no population, a difference between three documents rather than an inference about who may be treated. The oxybate row does not call itself an **orexin receptor antagonist** as the other three do; its label calls it a central nervous system depressant. And the last column is named for the field it came from: whether any is a first marketing approval these records do not say.

The record cautions about itself: openFDA is an agency-operated mirror of Drugs@FDA and of the label archive rather than the application itself, and the indication wording is current label text as filed, which is not a statement about any other jurisdiction and must never be carried across one. The reading dates differ: the three orexin rows carry a status true as of 1 August 2026 while the oxybate row carries 3 August 2026, because it was repaired in between. Harmonising the newer date onto the others asserts readings nobody made.

5.3 Right citation, right date, wrong product

The oxybate row is the most instructive failure here, because nothing about it looked wrong. The cited application and the date were both correct and both belonged to Xyrem. The indication text was not: it was the text of LUMRYZ, a different application by a different sponsor, approved in 2023, whose label reads almost identically in structure at 349 characters against 351 and differs mainly in the product name. Comparing lengths would have passed it; checking the retrieved document's own application number against the one cited caught it.

The vulnerability is structural: three of these four return one application, so there is nothing to blend; sodium oxybate returns four. Adjudication had already established that Xyrem, Xywav and Lumryz are three separate applications and that Xywav is not even a single-entity sodium oxybate product; that finding could not reach the ledger because its regeneration was parked, so the defect stood for a day after being solved elsewhere. The row is checkable now because it records an identity test anyone can repeat.

5.4 The European Union: two authorised, two absences, and a cut that closed with a full stop

Two agents carry a European centralised authorisation and two do not. For the two that do, indication text comes from Annex I of the **Summary of Product Characteristics**, section 4.1 Therapeutic indications, read from the cached product-information document; the register and assessment-report page corroborate identity and status only. Daridorexant: QUVIVIQ is indicated for the treatment of adult patients with insomnia characterised by symptoms present for at least 3 months and considerable impact on daytime functioning. The oxybate: Treatment of narcolepsy with cataplexy in adult patients, adolescents and children from the age of 7 years. Both are Authorised, daridorexant under EMEA/H/C/005634 and the oxybate under EMEA/H/C/000593, each with a **marketing authorisation holder** the register names by company and not by country: Idorsia Pharmaceuticals Deutschland GmbH and UCB Pharma Ltd. Both dates are of a declared kind, the initial European marketing authorisation as a Commission decision: **29 April 2022** and 13 October 2005, each file carrying a later Commission decision in 2026 and 2025.

> **GOOD TO REMEMBER**
>
> Whether a patient must have cataplexy to fall inside the licence depends on which regulator you are reading: one text reaches excessive daytime sleepiness in narcolepsy without requiring cataplexy, the other names narcolepsy with cataplexy. When a colleague quotes an oxybate indication, the question is not what it says but which country's document it came from.

That row carries this section's instrument lesson. Until it was corrected that row read "Treatment of narcolepsy with cataplexy in adult patients." The words "adolescents and children from the age of 7 years" had been dropped and the surviving prefix closed with a full stop the source does not have. The same correction fixed a second defect: the indication had come from the published register while the assessment-report page was cited as an independent second instrument, though the register is generated from those pages, so the corroboration was circular and blind to the stale summary layer that had failed.

TRAP **Treating a sentence that ends in a full stop as one that ends where the source ends.** A truncation landing mid-word announces itself; one landing at a clause boundary and closing with valid punctuation is undetectable by eye, and it removed a paediatric population from an authorised indication in this ledger. The instrument that catches it asks one question: is the stored text a prefix of the primary text? Learn the question, not the incident.

For the two with no European record the absence is written as an absence, its strength stated. The agency's own medicines register, generated on 1 August 2026, was scanned whole across all 39 columns of all 2730 medicine rows, and each substance returns zero rows under any of its names. The register carries Authorised, Refused, Application withdrawn, Withdrawn, Lapsed, Expired, Revoked, Suspended and Opinion outcomes, so absence means no central application of any kind was ever concluded, including a refused or abandoned one; what would close the row is an agency assessment report or a Commission implementing decision. The authority a claim would rest on is named, and "approved internationally" is not a permitted substitute for a named regulator.

Then the limit: the register covers the European **centralised procedure** only, and a product authorised nationally through the decentralised or mutual-recognition route would never appear in it. That limit is printed identically on every European absence here but does not bite equally, and on both these agents it is graded as biting weakly, because a small molecule outside the mandatory centralised scope could hold a national authorisation this instrument cannot see. Two weaknesses travel with the finding: the confirming second instrument ran on all eight grounded rows and none of the eight absences, so no European absence here has independent corroboration, and the assessment-report page would anyway be a second rendering of the register's own source.

5.5 The ladder that does not exist

Lemborexant carries a located Indian approval record and a declared European absence. That defeats the assumption almost every candidate brings to a recent-advances question: that maturity runs America, then Europe, then India. Here the middle rung is the empty one, and these four agents run three patterns across three jurisdictions.

The Indian record is affirmative. The approval-list cell reads, whole, Treatment for insomnia. Lemborexant tablets 5mg/10mg. It sits at serial number 12 of a 2021 approval list, one document in a corpus of 66 distinct published approval documents, deduplicated by content hash, discovered by four mapping passes and downloaded directly, all 66 yielding an extractable text layer, so no absence in this corpus is an artefact of an unreadable file. The strengths there are a formulation fact from the regulator's own name column, quoted because the discipline is to carry a cell whole; they are not a dosing recommendation, and no dose or maximum appears here.

The date needs the same care. The value is **1 July 2021**, and the field says what is printed: the figure in the column headed **Date of issue**. An earlier version claimed to quote a label the document does not carry, and was narrowed. Corroboration is thinner than it looks: two documents are cited, but they are one running list in two renderings rather than two independent documents, and the second lies outside the 66-document corpus, found by a drug-specific search

and appearing in no mapping pass, so the corpus discovery step was incomplete though the fetching step was not.

TRAP **Answering at the grain of the agent instead of the agent and the country together.** "Lemborexant is approved" is not a shorter version of a true sentence but a different and unverifiable one. Treat each of these twelve records as its own row, and answer any availability question with the jurisdiction, the record and the date. The mirror-image error costs as much: inferring a European or American absence from an Indian gap, when here the Indian record is the affirmative one.

5.6 The three Indian absences, and the one built on documents that exist

The other three carry no located Indian marketing approval, and the absences are not equal. Suvorexant's is strongest, resting on documents that exist rather than on documents that do not. The row records the file at pre-approval stage: CDSCO subject expert committee minutes for neurology and psychiatry of 24 September 2025 record a named manufacturer proposing permission to manufacture and market suvorexant tablets, the committee recommending permission to conduct a bioequivalence study, and minutes of 26 November 2025 record a second manufacturer with the same proposal, the committee recommending a Phase III clinical trial. Both are applications, and an application is not an approval.

That row's lead sentence once read that the agent was not approved in India at retrieval, and it was withdrawn: pre-approval activity by two named firms is evidence that those two firms had not yet been granted permission, not proof that no other holder exists, in a corpus this build itself shows to be incomplete. The evidence did not change; only the claim it was asked to carry was narrowed, and the row remains a declared absence.

The other two rest on searches that returned nothing usable, each naming what it searched. For daridorexant the substance is absent from all 66 published lists; four committee minutes returned by a domain search were all fetched, are all readable and none names the drug, and a fifth hit in an approval-list directory does not contain the string either, so all five were search-index false positives; no indexed CTRI trial record was located. For the oxybate the five domain hits are all International Cell written confirmation files naming Indian manufacturers, export-related documents that are not Indian marketing authorisation; no committee minute, no clinical-trial approval and no indexed CTRI record was located. A **written confirmation** is likeliest to be mistaken for an approval, being a real regulatory document about a real manufacturer that says nothing about the Indian market.

The authority that would settle any of these is the Central Drugs Standard Control Organisation and the Drugs Controller General of India, and committee minutes, clinical-trial no-objection certificates, International Cell written confirmations, CTRI registration, brand listings, pharmacy availability and manufacturer copy are none of them marketing authorisation. A **CTRI registration** records that a trial exists, not permission to market. Two sensitivity facts then bind any Indian sentence here. Every one of the seventeen positive controls this corpus finds sits in a document dated 2023 or earlier, and measured independent sensitivity in the 2024 to 2026 window is 0 of 1, so the corpus has never been shown to find a recent approval; and the positive-

control tally of 17 of 20 may not be cited as warrant for an absence on any agent approved after 2023. Absence from this corpus is evidence and not proof, and no Indian row in this build asserts that a drug is not approved in India: each asserts that no CDSCO approval record was located, and names what was searched.

IN INDIA

Write it one agent at a time. Lemborexant: a CDSCO approval-list entry for insomnia at serial number 12 of a 2021 list, date of issue 1 July 2021. Suvorexant: no CDSCO marketing approval located, the file documented at pre-approval stage on 2025 committee minutes; an application is not an approval. Daridorexant and sodium oxybate: no CDSCO approval record located, and what was searched is named. Do not write that any of them is not approved in India, and do not write that any is approved without naming the located record.

5.7 What certified these rows, and what a null cell is worth

These twelve records were tested rather than trusted, on three instruments: is the stored text verbatim against the primary document, is it that text cut short, does the count of indications carried match the document's? That is 36 cells, of which 21 are applicable and all 21 pass, with no failures and nothing unresolved.

1 of 2730	21 of 21	349 of 349	1 of 66
EU REGISTER MEDICINES NAMING OXYBATE	APPLICABLE FALSIFICATION CELLS PASSING	CHARACTERS, US OXYBATE ROW VERSUS LABEL	CDSCO LISTS WITH A LOCATED ENTRY

The fifteen inapplicable cells matter as much, and the reason for skipping them defines a declared gap: a row that is a declared absence carries no indication text, so an instrument comparing text against a primary document has nothing to test, because a gap is not a claim about a product but a claim about a search. An absence can be certified only as honest about its own reach.

Two instruments were proved on defects taken from this group, not invented ones. The prefix instrument was tested by replanting the real European oxybate truncation and returned a failure; the verbatim instrument by replanting the real Indian misattribution, in which a flattened text extraction attached the obesity indication belonging to the next serial number, cetilistat, to the lemborexant row, and it also returned a failure. That defect was repaired by reading the raw column layout of the source table instead of the flattened text line. An instrument never shown to fire on a known defect is not evidence.

The Indian cell is certified on a geometric reader's agreement with the drawn rules, its locator recorded to document, page, serial number and column coordinates, while the second document that row cites is absent from the cached corpus altogether. Certification concerns the reading, corroboration the sources.

5.8 The four records side by side, and what a complete answer contains

Reproduce this grid and the item is answerable in any form.

AGENT	UNITED STATES	EUROPEAN UNION	INDIA
Suvorexant	Located label indication for insomnia, NDA204569	No centralised procedure of any kind on record; scope limit weak	No approval located; file at pre-approval stage on 2025 minutes
Lemborexant	Located label indication for adult insomnia, NDA212028	No centralised procedure of any kind on record; scope limit weak	Approval-list entry for insomnia, serial 12, date of issue 1 July 2021
Daridorexant	Located label indication for adult insomnia, NDA214985	Authorised, EMEA/H/C/005634; indication carries a duration and a functioning criterion	No approval located; four minutes read, none names the drug
Sodium oxybate	Located indication: cataplexy or excessive daytime sleepiness in narcolepsy from age seven	Authorised, EMEA/H/C/000593; indication is narcolepsy with cataplexy including children from age seven	No approval located; the five documents found are export written confirmations

Reduced to four sentences:

- Status belongs to an agent inside a jurisdiction on a date, so twelve records is the right number and one class verdict is wrong.
- The two families carry different diseases in their licences, and their indications are never substituted for one another.
- The commonest silent defect is a regulator's sentence cut at a clause boundary and closed with a full stop.
- A declared absence is a statement about a search, naming the instrument, its reach, and whose authority would make it a negative.

MNEMONIC

Agent, Border, Cell, Cut

Agent: one drug in one country, never a class. Border: no indication, date or verdict crosses a jurisdiction. Cell: quote the regulator's cell whole, including the population at its end. Cut: ask whether a quoted indication is the primary text closed off tidily.

The licence belongs to a drug in a country on a date, and a sentence that ends in a full stop has not thereby ended where its source did.

6 · Xanomeline-trospium and muscarinic agonist antipsychotics

One capsule holds a muscarinic agonist and a muscarinic antagonist, and that pairing is the design rather than an accident of formulation. Two disciplines are examined at once. The

pharmacological question is how a treatment for psychosis that gives up dopamine blockade is supposed to work, and why the second molecule is there to be kept out of the brain rather than let into it. The regulatory question is why three national records of the same drug read so differently. Marks are lost in three predictable places: candidates treat the antagonist as a central anticholinergic; they carry an American indication across a border; and they read the silence of an instrument never shown able to speak. Receptor principles for the muscarinic system belong to the Psychopharmacology I: principles and core drug classes notes and are not rebuilt here.

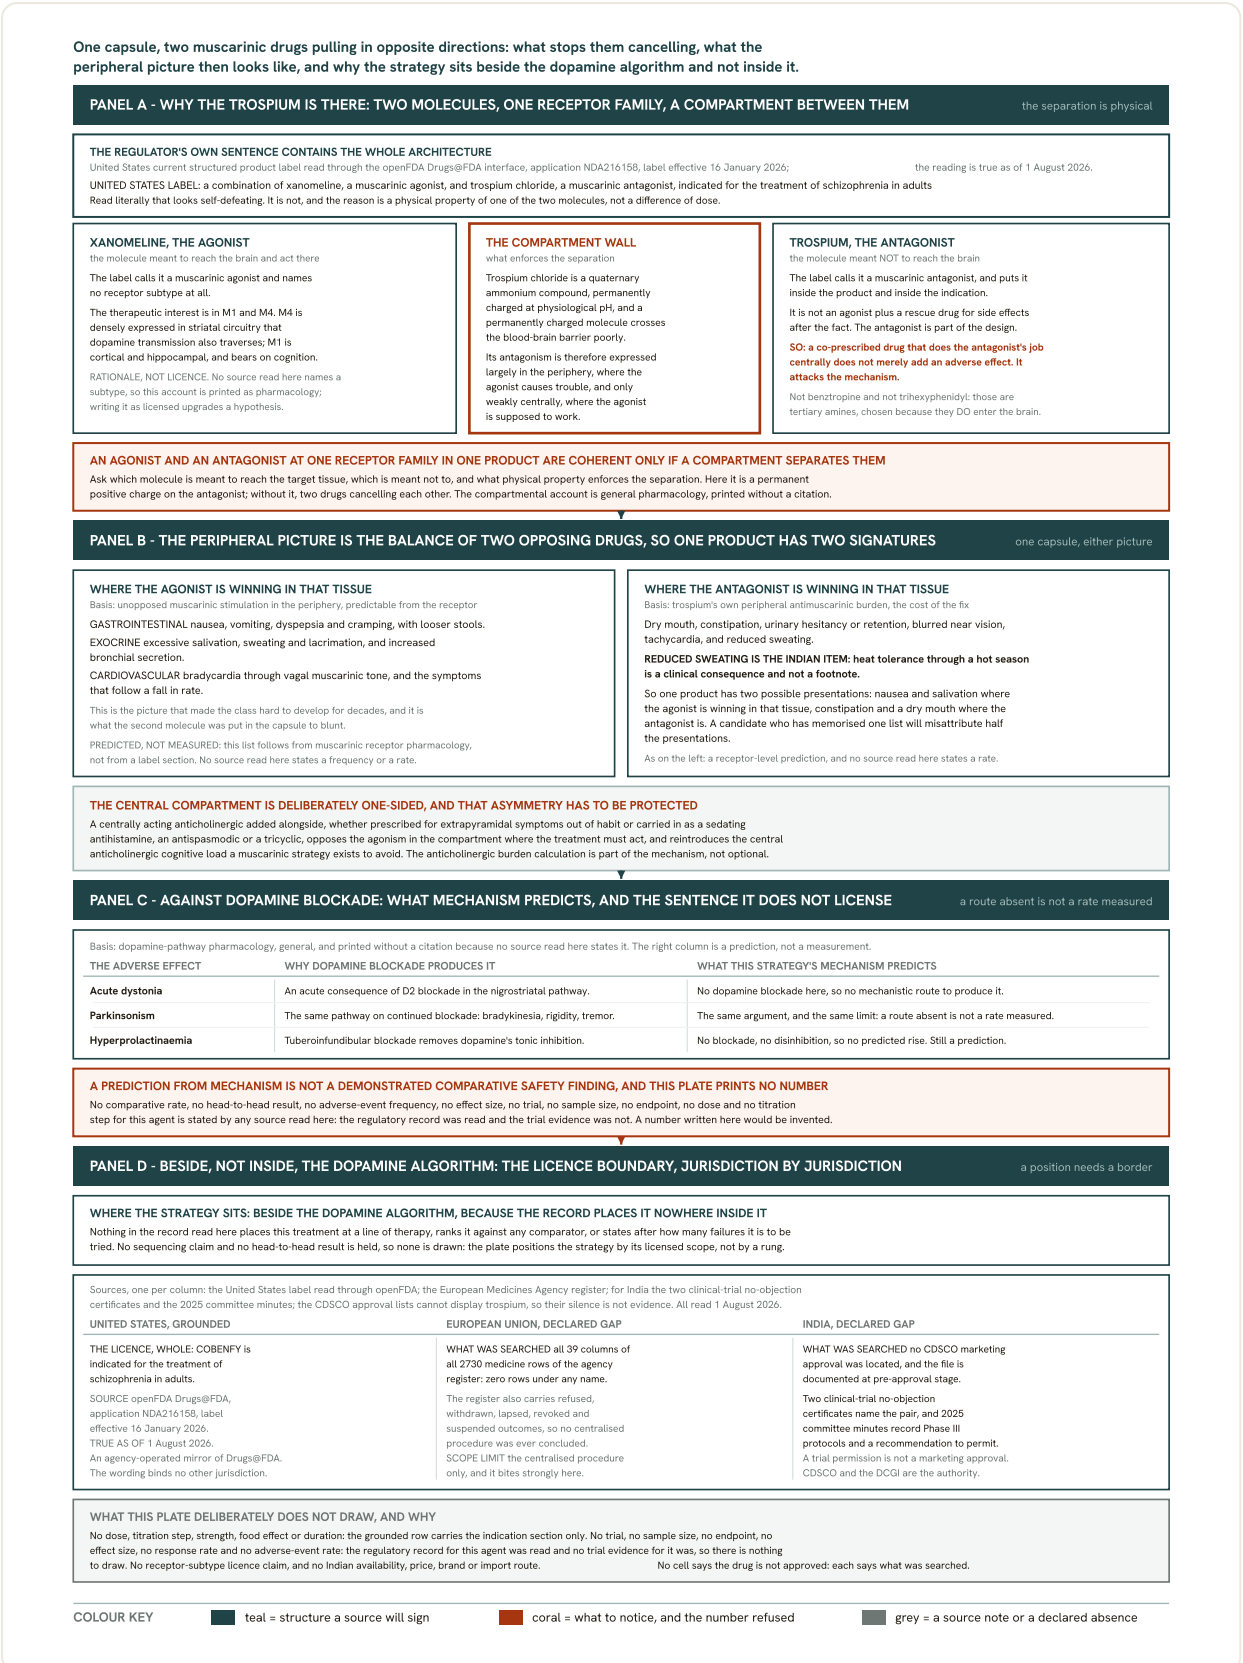


Fig 38.4 Muscarinic agonism with a peripherally restricted antagonist: why the trospium is in the capsule, and where the strategy sits. The label describes a combination of a muscarinic agonist and a muscarinic antagonist, and the pairing is coherent because trospium is permanently charged and stays peripheral, so one product can present with either signature, both predicted from the receptor and neither with a rate. The licence is the treatment of schizophrenia in adults in one jurisdiction; the European Union and India are drawn as declared gaps naming what was searched. (United States label via openFDA, NDA216158, effective 16 January 2026; European Medicines Agency medicines register; for India two trial no-objection certificates and 2025 committee minutes; read 1 August 2026.)

6.1 An agonist and an antagonist in one capsule, and why that is not a contradiction

Start with the sentence the regulator wrote; it contains the whole architecture. The product is described on its current United States label as a combination of xanomeline, a muscarinic agonist, and trospium chloride, a muscarinic antagonist, indicated for the treatment of schizophrenia in adults. Read literally that looks self-defeating; a candidate who says so has noticed the right thing and drawn the wrong conclusion.

The pairing is coherent because the two molecules are separated by a compartment and not by a dose. Xanomeline, the **muscarinic agonist**, is intended to reach the brain and act there. Trospium, the **muscarinic antagonist**, is a quaternary ammonium compound, permanently charged at physiological pH, and a permanently charged molecule crosses the blood-brain barrier poorly. Its antagonism is therefore expressed largely in the periphery, where the agonist causes trouble, and only weakly centrally, where the agonist is supposed to work.

Two things follow. The combination is not an agonist plus a rescue drug for side effects after the fact; the antagonist is inside the product and inside the indication. And a co-prescribed drug that does the antagonist's job centrally does not merely add an adverse effect. It attacks the mechanism.

RULE **An agonist and an antagonist at the same receptor family in one product are only coherent if a compartment separates them.** Ask which molecule is meant to reach the target tissue, which is meant not to, and what physical property enforces the separation. Here it is a permanent positive charge on the antagonist; without it, two drugs cancelling each other.

6.2 What the muscarinic strategy replaces, and where it sits

Every antipsychotic algorithm a postgraduate has learned is built on occupancy of the dopamine D2 receptor. This strategy is not on that axis; it engages the cholinergic muscarinic system instead, and the therapeutic interest sits with the M1 and M4 subtypes, M4 because it is densely expressed in striatal circuitry dopamine transmission also traverses, M1 because of its cortical and hippocampal distribution and its relationship to cognition. Xanomeline came out of a development programme aimed at Alzheimer's disease, where the antipsychotic signal was noticed before it was pursued in schizophrenia; the cholinergic tolerability problem that limited it then is what the second molecule was added to solve.

Be exact about what is licensed and what is merely explanatory. The label text held here calls xanomeline a muscarinic agonist and names no receptor subtype at all. The M1 and M4 account is pharmacology and belongs to the rationale, not the licence. Teach it as the rationale; writing it as though the regulator had endorsed a subtype-selective claim upgrades a hypothesis into a licence.

The strategy sits beside the dopamine-blockade algorithm rather than inside it, because nothing in the record read here places it at a line of therapy, ranks it against any comparator, or states after how many failures it is to be tried: no comparative efficacy claim, no sequencing claim and no head-to-head result is held for this agent in this build, and inventing one is more costly than admitting the absence.

GOOD TO REMEMBER

Acute dystonia, parkinsonism and hyperprolactinaemia are consequences of dopamine blockade, so a treatment that does not block dopamine has no mechanistic route to produce them. That is a prediction from mechanism rather than a demonstrated comparative safety finding, and what cannot be said without a trial is any figure comparing rates. Write the mechanism and stop; do not smuggle in a number.

6.3 The peripheral antagonist: what it buys and what it costs

An unopposed muscarinic agonist in the periphery produces a recognisable picture, the reason this class was hard to develop for decades. Predictable from the receptor alone:

- Gastrointestinal stimulation: nausea, vomiting, dyspepsia and cramping, with looser stools.
- Exocrine stimulation: excessive salivation, sweating and lacrimation, and increased bronchial secretion.
- Cardiovascular slowing through vagal muscarinic tone: bradycardia and the symptoms that follow a fall in rate.

Trospium is in the capsule to blunt exactly that list. The cost is that it brings its own peripheral antimuscarinic burden: dry mouth, constipation, urinary hesitancy or retention, blurred near vision, tachycardia, and reduced sweating with its implications in Indian summer heat. The net peripheral picture is therefore a balance of two opposing drugs given together, and the same product can present with either signature: nausea and salivation where the agonist is winning in that tissue, constipation and a dry mouth where the antagonist is. A candidate who has memorised one list will misattribute half the presentations.

The central compartment is deliberately left one-sided, and that asymmetry must be protected. A centrally acting anticholinergic added alongside, whether prescribed for extrapyramidal symptoms out of habit or carried in as a sedating antihistamine, an antispasmodic or a tricyclic, opposes the agonism in the compartment where the treatment must act, and reintroduces the central anticholinergic cognitive load a muscarinic strategy exists to avoid. The anticholinergic burden calculation is not optional here; it is part of the mechanism.

PITFALL

Describing trospium as though it were benztropine or trihexyphenidyl. Those are tertiary amines chosen because they enter the brain, to oppose central cholinergic tone. Trospium is chosen for the opposite property. Calling it an anticholinergic without saying peripherally restricted loses the point of the combination, and is the single sentence most likely to lose the mechanism mark here. The related error is prescribing a central anticholinergic alongside the product.

6.4 The United States: the one grounded record, read field by field

Of the three jurisdictions, exactly one carries a grounded status row, a model of a status claim written properly. The indication, verbatim: COBENFY is indicated for the treatment of schizophrenia in adults. The source is named and narrow: the status was read from the openFDA

Drugs@FDA interface, application NDA216158, with the current structured product label effective 16 January 2026, and the reading was true as of 1 August 2026. Jurisdiction, indication in the source's words, named source, date: the sentence becomes checkable by someone who does not trust you.

Then the field easiest to misreport. The earliest approval-status date in that application record is **26 September 2024**, and the field is named for what it holds, not for what a reader would like it to mean. Whether that date is a first marketing approval, a supplement or something else is a question the record read here does not answer, and this ledger has already published one drug's generic approval date under a first-approval field name and had to retract it. Label a date with the name of the field it came from, every time.

The record carries a caution about itself, which a careful answer respects: openFDA is an agency-operated mirror of Drugs@FDA and of the label archive rather than the Drugs@FDA application itself, and the indication wording is current label text as filed, which is not a statement about any other jurisdiction and must never be carried across one.

This row was tested rather than trusted. The query returned exactly one application, NDA216158, so the row does not blend several applications into one agent-level claim; and three text instruments were run against the primary document, checking that the row text is verbatim, that it is not the primary text cut short and closed with a full stop so that nothing signals the loss, and that the count of indications carried matches the count in the document. All three passed: 262 characters on the row against 262 in the primary, and one indication in the document against one on the row.

1 of 1	0 of 2730	0 of 66	0 of 121
INDICATIONS CARRIED, US ROW VERSUS LABEL	EU REGISTER MEDICINE ROWS MATCHING	CDSCO APPROVAL LISTS WITH AN ENTRY	EXTRACTED TEXTS NAMING TROSPIUM

TRAP **Widening the indication because the mechanism sounds as though it should be wider.** The licensed wording read here is **the treatment of schizophrenia in adults** and nothing else. It is not an indication for negative symptoms as a separately licensed target, for cognitive impairment, for psychosis in dementia, for adolescents, or for use as an add-on to another antipsychotic. The M1 cognitive rationale in 6.2 is what tempts a candidate to widen the sentence, and the widening is unlicensed. Quote the indication whole and let the mechanism stay a mechanism.

6.5 The European Union: an absence in a register that records refusals too

The European position is a declared gap, and a gap's strength depends entirely on the instrument that produced it. The instrument: the European Medicines Agency's own medicines register, generated by the agency on 1 August 2026, downloaded whole and scanned across all 39 columns of all 2730 medicine rows; this substance returns zero rows under any of its names.

What makes that more than a shrug is the range of outcomes the register carries. It records Authorised, Refused, Application withdrawn, Withdrawn, Lapsed, Expired, Revoked, Suspended and Opinion outcomes, so the absence means no central application of any kind was ever

concluded, including one refused or abandoned. The missing document is an agency public assessment report for the substance, or a European Commission implementing decision. Nothing else closes it, and the phrase approved internationally is not a permitted substitute for a named regulator.

Now the limit, where an honest answer separates itself. This register covers the European Union **centralised procedure** only. A product authorised nationally in one or more member states through the decentralised or mutual-recognition procedure would never appear in it. That limit is stated on every European gap row in this build, but it does not bite equally on all of them, and on this agent it is graded as biting strongly: this is a small molecule outside the mandatory centralised scope, so a national authorisation somewhere in the Union is genuinely possible and is wholly unmeasured here. Two further weaknesses are recorded and travel with the finding: the confirming second instrument was run on the grounded rows and on none of the gap rows, so no European gap in this build has independent corroboration, and the assessment-report page would in any case be a second rendering of the same source rather than an independent one, since the register is generated from those pages.

The defensible European sentence is therefore narrow, and worth writing in full: no European Union centralised procedure of any kind is on record for this active substance as of 1 August 2026, which does not exclude a national authorisation that this instrument cannot see.

6.6 India: the affirmative documents, and the sentence that was corrected away

The Indian entry is the one an Indian examination cares about, and it is the one where a confident wrong sentence is easiest to write. The row states that no CDSCO marketing approval was located and that the file is documented at pre-approval stage. That second clause is worth learning, because unlike most gaps in this chapter it is built on documents that exist rather than on documents that do not.

What was found is specific and dated. Two CDSCO **clinical-trial no-objection certificates** name Trospium Chloride plus Xanomeline, one against an IQVIA file numbered 44439 and one against a Bristol-Myers Squibb file numbered 52364. Minutes of the **subject expert committee** for neurology and psychiatry dated 30 May 2025, 10 December 2025 and 22 December 2025 record Bristol-Myers Squibb India Private Limited presenting the Phase III protocols CN0120051 and CN0120025, with the committee recommending permission to conduct the trial.

That footprint licenses one sentence and forbids another. A trial permission is not a **marketing approval**. The authority that would settle the question is the Central Drugs Standard Control Organisation and the Drugs Controller General of India, and expert committee minutes, clinical-trial no-objection certificates, International Cell written confirmations, CTRI registration, brand listings, pharmacy availability and manufacturer copy are none of them marketing authorisation. A **CTRI registration** deserves separate mention because it is the substitution an Indian candidate is likeliest to make: it records that a trial exists, not permission to market. What would

close the row is an affirmative CDSCO or DCGI new-drug permission, a current CDSCO-approved label, or an entry in an official CDSCO approval list.

Note what the row does not say, because it was deliberately taken out. The lead sentence of this row once read that the agent was not approved in India at retrieval, and that was corrected: pre-approval activity by two named firms is evidence that those firms had not yet been granted permission, and it is not proof that no other holder exists, in a corpus this build itself demonstrates to be incomplete. The evidence did not change; only the claim it was being made to carry was narrowed, and the row remains a declared gap either way. No India row in this build asserts that a drug is not approved in India. Each asserts that no CDSCO approval record was located, and names what was searched.

IN INDIA

Write it in this order. No CDSCO marketing approval record for xanomeline with trospium was located as of 1 August 2026; the Indian file is documented at clinical-trial stage, with two trial no-objection certificates and subject expert committee minutes of 2025 recording Phase III protocols and a recommendation to permit the trial; a trial permission is not a marketing approval. Do not write that the drug is not approved in India, which is stronger than this evidence supports. Do not write that it is approved, because nothing retrieved says so. For a patient or a family: the treatment is not something an Indian service can offer today outside a trial, and a sponsor's trial permission is not an availability promise to a patient.

6.7 The instrument that cannot see half of this drug

The India row carries one clause that looks like evidence and is worth nothing, and why is the most transferable thing on this item: the agent is absent from all 66 published CDSCO approval lists.

Trospium is absent from all 66 of those documents and from all 121 extracted text files under every spelling tried, including variants tolerant of the rn and m ligature confusion, of 0 against o and 1 against l, and of letter-spaced layout, and under the brand and programme names. It is one of three drugs that failed as **positive controls** for this corpus, alongside agomelatine and molnupiravir: drugs taken to be approved and marketed in India that the corpus nonetheless cannot display. The corpus has been shown incapable of displaying trospium at all, so its silence is a property of the instrument and not a fact about the drug.

The row records this explicitly: the absent-from-66-lists clause is worth nothing for trospium specifically and must not be read as a negative finding about it, and the verdict does not depend on that clause, because the verdict rests on the affirmative documents, the two trial no-objection certificates and the dated committee minutes, which are unaffected. The template: a null reading from an instrument shown blind to the target contributes nothing, and the same page can still carry a sound conclusion resting on something else.

Two sensitivity facts belong to any Indian gap sentence on a recent agent. Every drug this corpus is known to find sits in a document dated 2023 or earlier, and independent sensitivity in the 2024 to 2026 window is 0 of 1, so the corpus has never been shown to find a recent approval;

and the positive-control tally of 17 of 20 may not be cited as warrant for a gap on any agent approved after 2023. The corollary is that a gap is a statement about a search and not about a product, which is also why the three text-comparison instruments were recorded as not applicable to both gap rows on this agent: a declared gap carries no indication text, so an instrument that compares text against a primary document has nothing to test.

TRAP **Quoting a null result from an instrument that was never shown able to detect the thing.** Placed beside two genuine documents, the sentence that trospium appears in none of 66 approval lists reads as a third piece of evidence, and is not evidence at all. The test before quoting any absence: was this instrument ever demonstrated to return a known positive? Here the answer is no, recorded as one of three control failures. The mirror-image error is to conclude the whole row collapses because one clause is empty; the verdict was built on the affirmative documents and never on the absence.

6.8 The three records side by side, and what a complete answer contains

The same drug, the same date, three records of entirely different kinds. Reproduce the differences and the question is answered in any form.

JURISDICTION	WHAT THE RECORD HOLDS	THE SENTENCE THAT MAY BE WRITTEN	THE SENTENCE THAT MAY NOT
United States	A current label indication for the treatment of schizophrenia in adults, application NDA216158, read 1 August 2026	The indication quoted whole, with the application, label date and reading date	Any wider indication, and any date labelled as a first approval
European Union	Zero rows in the agency medicines register of 2730 medicines, which also records refusals and withdrawals	No centralised procedure of any kind is on record, as of the reading date	Not authorised in Europe, which the centralised-only instrument cannot support
India	Trial no-objection certificates and 2025 committee minutes recording Phase III protocols and a recommendation to permit, and no located approval	No CDSCO approval record located; the file is documented at trial stage	Not approved in India, and any negative inference from the approval-list corpus about trospium

Reduced to four sentences, the item is this.

- The mechanism is muscarinic, not dopaminergic; the combination works because a permanently charged antagonist stays peripheral while the agonist acts centrally.
- The adverse-effect picture balances two opposing peripheral actions, so either a cholinergic or an antimuscarinic presentation is possible from one product, and a centrally acting anticholinergic added alongside attacks the mechanism.
- The regulatory position has three answers and no single one, each carrying its jurisdiction, its indication in the source's words, its named source and the date it was read.

- The Indian file is a declared gap standing on affirmative documents, and its one clause reporting an absence is void because the instrument is blind to trospium.

MNEMONIC

Charge, Compartment, Country, Control

Charge is why trospium stays out of the brain. Compartment is where each molecule is meant to act, and why the pairing is not self-cancelling. Country is the jurisdiction every status sentence must name and never leave. Control: before quoting any absence, ask whether the instrument was ever shown able to find what it was looking for.

The antagonist is in the capsule to be kept out of the brain, the approval is in one country and must stay there, and an absence reported by a blind instrument is not a finding.

7 · Endoxifen and protein-kinase-C-targeted agents in mania

This is the one agent in the chapter with an affirmative Indian record and a measured American absence, and the examinable content of the item is that both sentences are true at once. The marks sit on the regulatory question: writing about one molecule in three jurisdictions without letting any of them leak into another. It goes four ways. Writing "approved" with no country attached. Converting an American search that found nothing into a statement that nothing exists. Reading a tablet strength printed inside a product-name cell as a dose. And quoting the Indian indication with its qualifier quietly missing, the defect this build planted on this entry to prove its instruments could catch it. The acute-mania treatment algorithm is not rebuilt here.

7.1 A metabolite promoted to a drug, and what that changes

Endoxifen is an **active metabolite** of tamoxifen, and tamoxifen is named here only as the parent compound. When a parent owes its activity to a metabolite, response depends on a conversion step, and that step is where people differ: genotype, competing inhibitors and adherence all act on it before any target is reached. Give the metabolite itself and the step leaves the chain, taking its variability with it.

That is an argument about pharmacological architecture and not about magnitude. It says where variability enters, not how much of it there is, and no exposure figure, potency ratio, conversion proportion or genotype frequency is held anywhere in the material behind this chapter. It is also a different drug from its parent for regulatory purposes, with separate applications, indications and approval records, so nothing true of tamoxifen may be carried onto endoxifen in either direction.

GOOD TO REMEMBER

Giving a metabolite removes a source of variation between patients. It does not remove the need to show that the metabolite works. Removing a conversion step is a pharmacokinetic argument and the prescriber's question is a pharmacodynamic one, so answering the second with the first answers the wrong question.

7.2 The protein kinase C rationale, and what kind of claim a rationale is

Protein kinase C is not a receptor. It is a family of intracellular signalling enzymes downstream of receptor activation, acting on the machinery that changes what a cell does over hours and days rather than milliseconds. The intervention is aimed at a node inside the cell, several steps from the synapse, and the phenotype it must move is a whole manic episode.

The shape of that proposal is what an examiner can test. It is target-first: start from a mechanistic node, find a compound with activity there, and predict a clinical effect in a syndrome the node is thought to join. Most older agents arrived the other way round, from a clinical observation worked backwards. So its strength and its weakness are one feature: it is derived rather than observed.

A rationale licenses a hypothesis and a trial. It does not license a statement about efficacy, about how quickly an effect appears, about how large it is, or about how the agent stands against any existing treatment. No artefact behind this chapter holds trial evidence for this agent: no trial name, no sample size, no endpoint, no effect size, no response rate, no dose, no titration step, no adverse-event rate. This account names none of them, and offers no comparison against any existing mood stabiliser, because none is held.

RULE **A mechanistic rationale is a reason to run a trial and never a substitute for one.** Could your sentence be false even if the mechanism is right? "Protein kinase C is implicated in mania" and "inhibiting protein kinase C treats mania" are separated by the whole of clinical trial evidence.

7.3 India: the affirmative record, read field by field

The Indian record is the strongest of the three here, which reverses the order candidates expect. Its source is the published new-drug approval lists of the Central Drugs Standard Control Organisation, CDSCO: 66 distinct documents deduplicated by content hash, downloaded direct over verified transport on 1 August 2026, every one yielding an extractable text layer so that no absence in the corpus is an unreadable-file artefact; the entry is serial number 20 in the list file ndedec13appro.pdf.

The indication, quoted whole because a paraphrase is where the qualifier goes missing: For the acute treatment of manic episodes with or without mixed features of Bipolar I disorder. Endoxifen citrate bulk and Endoxifen tablets 8 mg. The first leg is an indication bounded three ways: acute treatment, manic episodes, bipolar I. The words **mixed features** are the part most often dropped, and dropping them changes who the indication covers. The second is a product-name leg naming a **bulk drug** entry and a tablet strength. That strength is a formulation fact

printed inside a name cell: not a dose, not a starting dose, not a maximum, and no dosing information for this agent exists in any artefact behind this chapter.

Then the date, the field this build has already had to correct once. The entry carries **10 December 2019**, and the row now describes that as the quantity printed in the column headed **Date of issue** rather than as an approval date, the earlier wording having claimed to quote a label the document does not print. The reading is dated too: the status true as of date on this row is 1 August 2026.

Corroboration is thinner than the verbatim quotation looks. It rests on one CDSCO document. Date and indication were confirmed against the raw column layout of the source table rather than its flattened text, and the neighbouring date of 20 November 2019 was checked and belongs to serial number 19 and to an entirely different substance. In a printed table a date can be read off the row above with nothing looking wrong.

An older record of the same search read this jurisdiction as a gap and named what would close it: a CDSCO or DCGI permission, an approved label, or an official approval-list entry with its date and indication verbatim; and what would not: a CTRI registration, a brand listing, pharmacy availability or manufacturer copy. What closed it was the third permitted item and nothing weaker. The question is never whether a document exists, but whether it can settle the question.

IN INDIA

Write it in this order. An entry for endoxifen appears in a CDSCO published new-drug approval list at serial number 20, carrying the indication above and the date 10 December 2019 under the printed heading Date of issue, true as of 1 August 2026. Quote the indication with its mixed-features qualifier intact. Do not turn the 8 mg tablet strength into a dosing instruction. And be clear what the record does not settle: it is an approval-list entry, not a statement about supply, price or whether your pharmacy stocks it.

7.4 How that Indian cell was certified, and the defect planted in it

Quoted correctly and read correctly are two achievements; this row was tested for the second. Three text instruments were run against the primary document: one checking that the row text is verbatim and that the source sentence names the primary, one checking that the row is not the primary cut short, at 146 characters against 146, and one checking that the indication cell is carried in full, the primary cell measuring 94 characters. All three passed against a locator naming the document, page 4, serial number 20, the bounds of the indication column and the row band, with corroboration recorded as trustworthy.

The rule behind that verdict: a cell counts as read when a reader working from the drawn rules of the table agrees with at least one independent reader, and two readers of the same flattened-text kind are not independent, because they share whatever made the first of them wrong.

Now the part that makes this row the most instructive object in the section. This entry is the one the build used to test its own instruments, planting two defects in it. The first cut the indication at a clause boundary and closed it with a full stop where the primary entry continues with or

without mixed features of Bipolar I disorder, dropping a qualifier that changes who the indication covers. The second removed the same qualifier from the middle of the cell and left the ending intact, so nothing was truncated at the end: that plant exists to separate the two instruments, because the prefix instrument must stay silent on it and only the completeness instrument may fire. Both guards fired and both plants were caught.

Note what both plants attacked. It was the clinical qualifier.

TRAP **Quoting an indication that ends in a full stop and assuming it ended there in the source.** Cut at its clause boundary, the Indian entry reads as a complete, plausible sentence with the mixed-features qualifier gone, and a cut closing with valid punctuation leaves no mark, so only comparison against the primary finds it. Reproducing an approved indication means reproducing its qualifiers: population, phase, severity, specifier. And a gate never shown to catch the defect it was built for has not been tested; planting the defect is the test.

3 of 3	0 of 5	228	0 of 2730
INDIA TEXT INSTRUMENTS PASSING	US QUERY ROUTES RETURNING A RECORD	LABEL DOCUMENTS MATCHING THE STRING	EU REGISTER MEDICINE ROWS MATCHING

7.5 The United States: a measured absence, and the sentence it does not license

An absence that was measured is a different object from an absence that was assumed, and only one of the two may be written down. A query of the Drugs@FDA interface returned no application record for the generic name endoxifen on 1 August 2026, the endpoint answering with its own not-found response. That probe found nothing; it does not establish that the agent is unapproved in the United States. To close the row a human must read Drugs@FDA directly and, if there is no application, say so on the agency's authority rather than on this probe's silence. The authority any American claim would rest on is the United States Food and Drug Administration, through Drugs@FDA.

The gap was then strengthened into a **measured absence**. Five query routes were run, every one returned the endpoint's own not-found response, there were zero transport failures, and every route was positive-controlled in the same run against a value known to return records, so a zero here is a route that answered and found nothing, not a route that failed. That **positive control** is the difference: without one, a broken route and an empty answer look identical from outside. The five routes were four fields of the application interface, generic name, substance name, active ingredient name and brand name, and one field of the label archive.

The scope is narrower than it first reads. It measures the Drugs@FDA interface and the label archive only. It is not the agency's own affirmative statement that no application exists, and it is silent on investigational and breakthrough designations, which appear in neither instrument. So the defensible American sentence names the search and stops. Writing it harder is forbidden, because an American gap in this build has already been wrong once: one agent's record carries no cross-reference block at all, so three of five routes returned not-found on it while the active-ingredient route returned the record.

Then the figure showing how a search can look conclusive in the wrong direction. A free-text query of the label archive for the string endoxifen, run during independent adjudication on 2 August 2026, returned a successful response with 228 matching documents, while every structured route above returned zero. The two are not in conflict. A **free-text search** matches a string anywhere in a document, so a document belonging to another product that mentions the substance counts as a hit, whereas a structured field query asks whether the substance is the product's own named ingredient.

TRAP **Collapsing an Indian approval-list entry and an American gap into one sentence about the drug.** This agent is the border-crossing error waiting to happen, both ways. Writing that endoxifen is approved, with no country attached, exports an Indian entry to the United States and Europe. Writing that it is not approved in the United States converts a search that found nothing into a regulator's negative finding, which no instrument here supports. Three jurisdictions, three records, three sentences, each with its own source and date.

7.6 The European Union: a register that records refusals, and the limit that follows

The European position is also a declared gap, and a gap is only as strong as the instrument behind it. The instrument is the European Medicines Agency's own published medicines register, generated by the agency on 1 August 2026, downloaded whole and scanned across all 39 columns of all 2730 medicine rows; this substance returns zero rows under any of its names.

What lifts that above a shrug is the range of outcomes the register carries. It records authorised, refused, application withdrawn, withdrawn, lapsed, expired, revoked, suspended and opinion outcomes, so an absence means no central application of any kind was ever concluded, including one refused or abandoned. The missing document is an agency public assessment report for the substance, or a European Commission implementing decision, and the phrase approved internationally is not a permitted substitute for a named regulator.

Now the limit. The register covers the European Union **centralised procedure** only, so a product authorised nationally through the decentralised or mutual-recognition procedure would never appear in it. That limit is printed identically on every European gap row, so it is graded per agent. For this substance it is graded as biting moderately: a small molecule outside the mandatory centralised scope could genuinely hold a national authorisation this instrument cannot see. The confirming instrument ran on all eight grounded European rows and none of the eight gap rows, so no European gap here has independent corroboration, and the assessment-report page would anyway be a second rendering of the same source. Second-instrument coverage of European gaps is **0 of 8**.

So the defensible European sentence is narrow: no European Union centralised procedure of any kind is on record for this substance as of 1 August 2026, which does not exclude a national authorisation this instrument cannot see. The three text instruments used on the Indian entry report nothing here, because a declared gap carries no indication text: a gap is not a claim about a product, it is a claim about a search. Each record was built in isolation, the Indian row consulting no American or European value, and the European row no American or Indian one.

7.7 The three records side by side, and what a complete answer contains

One molecule, one reading date, three records of different kinds.

JURISDICTION	WHAT THE RECORD HOLDS	THE SENTENCE THAT MAY BE WRITTEN	THE SENTENCE THAT MAY NOT
India	An approval-list entry at serial number 20 carrying the acute manic episode indication with or without mixed features, a bulk-drug leg and a tablet strength, dated 10 December 2019 under the heading Date of issue	The indication whole, with the document, serial number, date heading and reading date	Any dose built from the strength in the name cell, or any claim about supply
United States	Five positive-controlled query routes returning the endpoint's own not-found response, zero transport failures	These named routes, on this date, returned no application record	Not approved in the United States, which converts a probe's silence into the regulator's finding
European Union	Zero rows in an agency register of 2730 medicines that also records refusals and revocations	No centralised procedure of any kind is on record at the reading date	Not authorised in Europe, which a centralised-only instrument cannot support

Reduced to five sentences, the item is this.

- Three jurisdictions, three records, three sentences: an approval-list entry, a measured absence and a register absence, each with its own source and date.
- The Indian indication is quoted whole because the mixed-features qualifier decides who it covers, and the strength beside it is a formulation fact, not a dose.
- A protein kinase C target is a rationale for a trial, not a result from one, and no trial figure is held.
- An absence is a statement about a search: name the routes, their controls, and what the instrument cannot see.
- A date is only as strong as its field name, and a printed column heading is not a regulatory event.

MNEMONIC

Metabolite, Mechanism, Mixed features, Measured

Metabolite is what the drug is. Mechanism is a rationale until a trial says otherwise. Mixed features is the qualifier that must survive into your answer. Measured is the only kind of absence you may write down.

One molecule, three records: an Indian approval-list entry quoted whole with its qualifier, an American absence that was measured rather than assumed, and a European silence a national authorisation could still be hiding behind.

8 · Anti-amyloid monoclonal antibodies as a recent advance (cross-reference to dementia)

This item is a pointer, and knowing why it is a pointer is itself examinable. **Anti-amyloid monoclonal antibodies** reach a recent-advances list because they changed what a therapeutic claim here may sound like, and they sit in another chapter because their mechanism, trial evidence, imaging complication, genotype caveat, selection criteria, monitoring and regulatory record are taught with the Dementias and neurocognitive disorders notes. What is left is what this chapter owns: the class of claim these agents exemplify, the shape such an argument must take, why the regulatory position is not printed here, and what an Indian service must decide.

8.1 Why this is a pointer, and what a pointer owes the reader

This material is taught with the dementia notes, and the item itself is worded as a recent advance carrying a cross-reference there, so a pointer is what is being asked for and the reason is not space. Two accounts of one drug class in one paper are written at different times against different sources and drift apart, and the drift stays invisible until a candidate quotes the thinner one and is corrected out of the fuller one.

A pointer owes the destination named exactly, some idea of what is there, and a plain statement of what is withheld. The owning chapter carries dedicated passages on each of these agents and on the imaging complication, so the material is there and at length. Everything listed at the head of this section is taught with the Dementias and neurocognitive disorders notes and none of it is restated here. Going quiet on all of that without saying so invites a reader to fill the silence from memory.

8.2 The class of claim, and why psychiatry has an interest in it

Almost everything a psychiatrist prescribes makes a **symptomatic** claim: the current state will be better while the drug is taken, and it is judged on that state. A **disease-modifying** claim differs in kind: it asserts an effect on the course of the illness rather than on the state at a point along it, so it cannot be demonstrated at a point. It needs an observation period long enough for two trajectories to separate, an agreed way of locating a patient on one, and an argument that the gap between them matters to the person living it. Each is a place the claim can fail without the drug doing anything wrong.

Which is why the class belongs in a psychiatry chapter: these patients reach **old-age psychiatry** clinics as well as neurology ones, and the first clinician a family asks is often a psychiatrist. And the argument travels, because the way a disease-modifying claim is first pressed becomes the template the next class is held to.

RULE **A symptomatic claim is judged on the state, a disease-modifying claim on the course, and the same evidence cannot settle both.** Decide which claim is on trial before appraising any result: a better state at a point is silent on the course, and separating trajectories show nobody better today.

8.3 The shape of the evidence argument, with no figure from it

Such an argument runs at two levels and keeping them apart is where the marks are. **Target engagement** is the first: the thing the drug acts on has changed. Patient benefit is the second: the person does better for longer. Engagement alone is a **surrogate marker** argument, only as strong as the assumption joining marker to outcome, and the history of therapeutics is largely a history of that assumption failing. A marker result is a reason to keep testing, not to treat.

At the patient level the problem is what a difference on a rating scale means. A difference can be reliably present and too small for anyone to notice, and the bridge is the **minimally important difference**, the change below which no patient or family could tell. Where that threshold sits is contested rather than settled, so one honest appraiser reads a result as a worthwhile gain and another as clinically negligible; state both readings and name the threshold each rests on. Two constraints join the balance: a group difference in an average rate of change is not a prognosis for the individual in front of you, and where a treatment carries a delivery burden and a surveillance requirement the benefit side must weigh more before the balance tips.

No figure from that argument appears here, by rule and not oversight. The chapter holds no grounded evidence lane for this item: no trial name, no sample size, no observation period, no effect size, no adverse-event rate. Writing one from memory is a fabricated citation, the one defect that cannot be repaired later.

8.4 Why the regulatory position is not on this page, and what a delegated gap is

The regulatory position of these agents is not written here, and the reason earns more marks than the position would. The chapter's status ledger holds **48** rows and not one is an anti-amyloid row. That is neither an oversight nor a failed search: the ledger says no row is written there deliberately, so that none can be quoted from there, and it names the three agents covered, lecanemab, donanemab and aducanumab.

Three silences look identical here and mean different things. A searched gap is a question a named instrument was pointed at, which returned nothing; the honest sentence names what was searched and whose authority a claim would rest on, and never hardens into a categorical negative about a jurisdiction. An instrument blindness is a question the instrument could not have answered whatever the truth was, so its silence carries no information. A **delegated gap** is neither: the fact exists, another owner holds it, and this page declines to restate it so two accounts of one fact cannot diverge. This item is the third kind, and the only one that is not an absence of fact.

A delegation is not a warrant, though: nothing read for this chapter establishes the evidentiary state of the owner's account, which predates this chapter's four-field standard and carries its own warning against giving its record false precision. So go to the owner for what it holds, read

anything found there against the jurisdiction it was written about, and never carry a position across a border. Until that chapter is refreshed this class has a delegated status, not an unknown one.

0 of 48	3	131	16
CHAPTER STATUS ROWS ON THESE AGENTS	AGENTS WHOSE STATUS ANOTHER CHAPTER HOLDS	OWNER PASSAGES MATCHING LECANEMAB, DONANEMAB OR THE IMAGING COMPLICATION	OTHER AGENTS THIS CHAPTER HOLDS A STATUS FOR

TRAP **Supplying the regulatory position from memory because the section did not.** A reader given no jurisdiction, no wording and no date often finishes the sentence privately, and a half-remembered headline surfaces attached to the wrong country and year. A status sentence is worth writing only when it carries the jurisdiction, the wording the regulator used, the source it was read from and the date it was true; none of the four is available here for these agents.

8.5 What an Indian service is actually being asked to decide

The decision facing an Indian psychiatrist is not whether to prescribe one of these agents, and treating it as though it were is how the topic gets answered badly. Four decisions are live now, whatever any regulator has done. What to say to a family who arrive having read something: name what kind of treatment is discussed, say what is held elsewhere and can be checked, and offer to find out rather than guess, because a warm guess is repeated as though a doctor confirmed it. Whether raising the class unprompted is kind where any pathway would be long and uncertain, which is a judgment about this family and not a rule. What would have to exist before the question became a real option, which is the pathway set out below. And affordability, which in Indian practice usually settles the question first.

IN INDIA

Write the Indian answer in this order. The regulatory position is not stated here and is not to be inferred from the silence. Nothing read for this chapter holds an Indian position for these agents: the status ledger writes no row for them at all, and the material with the Dementias and neurocognitive disorders notes is written about the United States and Europe, its own gap list recording that current Indian regulatory, availability and cost information is not supplied there. So nothing found at the owner may be read across into an Indian answer. What an Indian old-age service must establish for itself is the pathway: who confirms the diagnosis, where the investigations are done, who owns the decision, who carries the follow-up, and what the family can afford.

8.6 The boundary in both directions, and what a complete answer contains

QUESTION	WHERE IT IS HELD	WHAT THIS PAGE MAY SAY
How the antibodies act, and on what	The owning item	Only that the class name states its target
What the trials showed, in numbers	The owning item	Nothing numerical: no evidence lane exists here
Imaging complication, genotype caveat, selection, monitoring	The owning item	Only that they exist and dominate the safety discussion
Regulatory position in any jurisdiction	The owning item, a delegated gap until refreshed	Nothing, for any agent, in any jurisdiction, on any date
The class of claim, and what a service must decide	Here	Claim type, argument shape, the delegation, the Indian pathway

Reduced to four sentences:

- A pointer by design: two accounts of one drug class diverge, and the thinner loses marks.
- The claim is disease-modifying, not symptomatic, which changes what evidence could settle it and how long it must run.
- The argument has two levels; a marker result licenses no patient conclusion, and the noticeability threshold is contested, so both readings are stated.
- The status is a delegated gap, not an unknown: held by its owner, restated nowhere here, and not yet meeting the four fields a status sentence needs.

GOOD TO REMEMBER

The commonest real-world error here is conversational, not pharmacological: a family told kindly and vaguely that something new exists and might help leaves having heard a promise nobody made. Name what is known, what is held elsewhere, and what you will find out.

MNEMONIC

Claim, Course, Custody, Clinic

Claim: symptomatic or disease-modifying, decided first. Course: judged over a trajectory, never at a point. Custody: another owner holds the record; this page restates none of it. Clinic: pathway, follow-up, affordability.

This item is a pointer: the class of claim and the service decision are taught here; mechanism, evidence and regulatory position are quoted from their owner or not at all.

9 · Moncrieff serotonin umbrella review and the serotonin-hypothesis debate

The examinable object here is a piece of evidence appraisal, not a drug and not a neurotransmitter. In 2022 a group of authors led by Joanna Moncrieff published an umbrella

review in Molecular Psychiatry asking whether the research literature supports the idea that depression is caused by low serotonin, reported that it does not, and set off an argument that reached patients before it reached examiners. Every mark in this topic is won or lost on discipline of attribution: what the review found, what its authors concluded, what its critics answered. The serotonin system itself, its pathways and its receptor subtypes, belongs to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes, and the monoamine hypothesis of depression and its limitations belong to the Mood disorders: classification, depressive and bipolar syndromes, neurobiology notes.

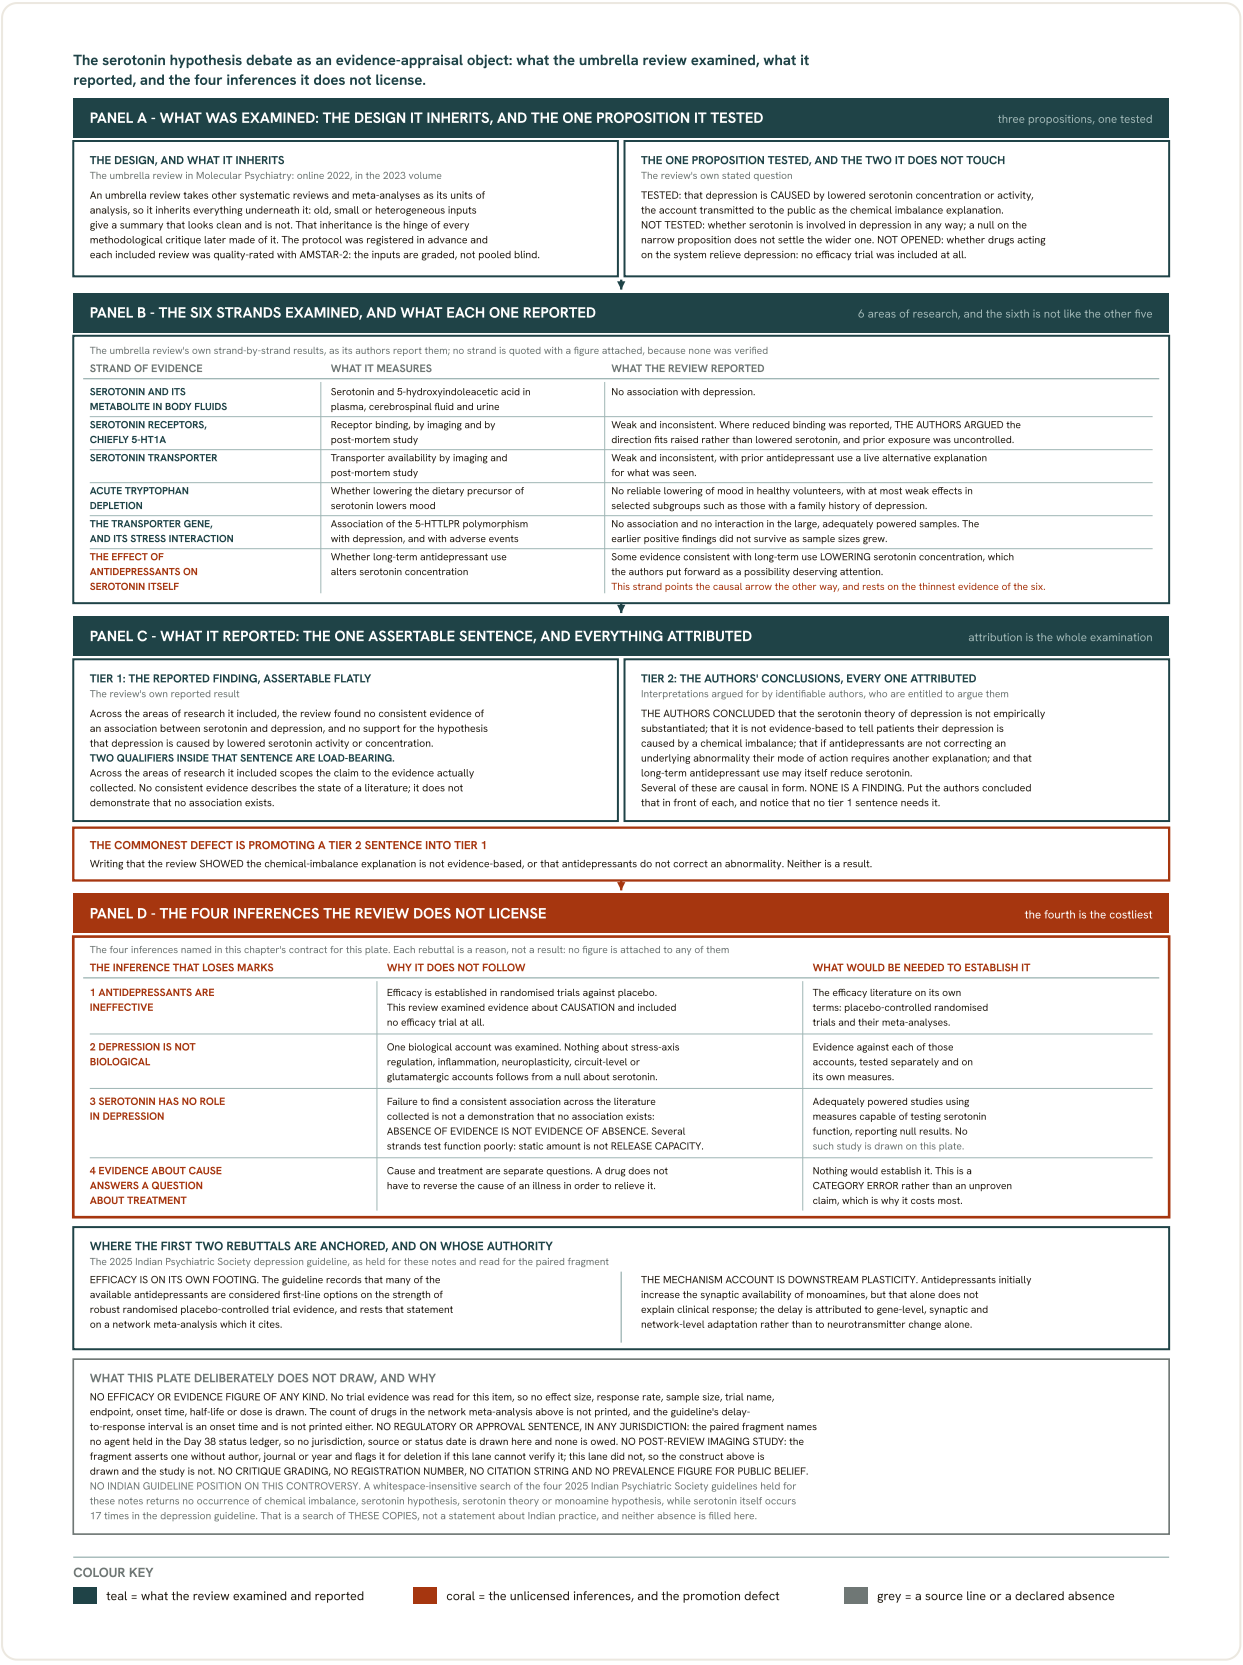


Fig 38.5 The serotonin hypothesis debate: what the umbrella review examined, what it reported, and the four inferences it does not license. One narrow proposition was tested, that depression is caused by lowered serotonin concentration or activity; neither the wider serotonin question nor efficacy was opened, because no efficacy trial was included. Six strands carry a reported finding and no figure, the sixth pointing the causal arrow the other way. What it reported divides into the finding, assertable flatly with both its qualifiers and the authors' conclusions, every one attributed. (The 2022 umbrella review in Molecular Psychiatry; the 2025 Indian Psychiatric Society depression guideline as held for these notes.)

9.1 What an umbrella review is, and why the design carries the whole argument

An **umbrella review** is a systematic review whose units of analysis are other systematic reviews and meta-analyses rather than primary studies. It sits at the top of the hierarchy of evidence and inherits everything underneath it. If the reviews it collects are old, small, heterogeneous or poorly conducted, so is the summary, however clean it looks. That inheritance is the hinge of every methodological criticism later made of this paper.

One housekeeping point before the method. The paper appeared online in 2022 and was carried in the journal's 2023 volume, so both years are correct and a marker who sees the other one has not caught an error.

Three features of the method matter for an answer. The protocol was registered in advance, which matters because a review of a contested hypothesis is exactly the kind whose scope can drift toward the answer its authors expect. The quality of each included review was rated with **AMSTAR-2**, a standard appraisal instrument for systematic reviews, so the paper does grade its own inputs rather than pooling them blind. And where no systematic review existed for a strand, the authors went to large individual studies instead. That decision is defensible, and also a departure from the pure design critics later used against the paper.

The review did not ask whether serotonin is involved in depression in any way. It asked whether the evidence supports a much narrower proposition: that depression is caused by lowered serotonin concentration or activity, the account transmitted to the public as the **chemical imbalance** explanation. Those are different propositions, and almost every error in this topic is a slide between them.

RULE **One narrow proposition was tested, and every sentence written about the review must name which proposition it is about.** "Depression is caused by low serotonin" is not the same claim as "serotonin is involved in depression", which is not the same claim as "drugs acting on the serotonin system relieve depression". The first was tested. The second was not tested and is not settled by a null result on the first. The third belongs to a different literature that this review did not open.

9.2 The six strands, and what each one reported

The evidence was organised into **six** areas of serotonin research, and the organisation is itself examinable: a question asking what the review examined is asking for this list. Reproduce it with a finding attached to each line, never as a bare list.

STRAND OF EVIDENCE	WHAT IT MEASURES	WHAT THE REVIEW REPORTED
Serotonin and its metabolite in body fluids	Concentrations of serotonin and of 5-hydroxyindoleacetic acid in plasma, cerebrospinal fluid and urine, compared between depressed and non-depressed groups	No association with depression
Serotonin receptors, chiefly the 5-HT1A receptor	Receptor binding, by imaging and by post-mortem study	Weak and inconsistent. Where reduced binding was reported, the authors argued the direction fits raised rather than lowered serotonin, and prior antidepressant exposure was not controlled for
Serotonin transporter	Transporter availability by imaging and post-mortem study, and its relation to depression	Weak and inconsistent, with prior antidepressant use a live alternative explanation for what was seen
Acute tryptophan depletion	Whether experimentally lowering the dietary precursor of serotonin lowers mood	No reliable lowering of mood in healthy volunteers, with at most weak effects in selected subgroups such as those with a family history of depression
The serotonin transporter gene, and its interaction with stress	Association of the 5-HTTLPR polymorphism with depression, and the proposed interaction between that polymorphism and adverse life events	No association and no interaction in the large, adequately powered samples; the earlier positive findings did not survive as sample sizes grew
The effect of antidepressants on serotonin itself	Whether long-term antidepressant use alters serotonin concentration	Some evidence consistent with long-term use lowering serotonin concentration, which the authors put forward as a possibility deserving attention

Read the sixth row differently from the first five. The first five ask whether a serotonin abnormality precedes depression; the sixth asks whether treatment produces one, which points the causal arrow the other way, and it is the strand carrying the authors' own hypothesis most visibly. It also rests on the thinnest evidence of the six. An answer that reports it with the same confidence as the genetic strand has flattened the paper into a slogan.

9.3 Three tiers: the finding, the authors' reading, and what neither warrants

Attribution is the whole examination here, so build the answer in three labelled tiers. The first tier is the review's reported finding, and it may be asserted flatly: across the areas of research it included, the review found no consistent evidence of an association between serotonin and depression, and no support for the hypothesis that depression is caused by lowered serotonin activity or concentration. Two qualifiers inside that sentence are load-bearing. "Across the areas

of research it included" scopes the claim to the evidence actually collected. "No consistent evidence" describes the state of a literature; it does not demonstrate that no association exists.

The second tier is what the authors concluded, and every sentence of it carries their name. They held that the serotonin theory of depression is not empirically substantiated; that it is not evidence-based to tell patients their depression is caused by a chemical imbalance; that if antidepressants are not correcting an underlying abnormality then their mode of action requires some other explanation; and that long-term antidepressant use may itself reduce serotonin. These are interpretations, several of them causal in form, argued for by identifiable authors who are entitled to argue for them. They are not findings, and a script that presents them as findings has surrendered the marks this topic exists to award.

The third tier is the set of inferences that neither the findings nor the authors' conclusions warrant. Four recur, and an examiner is watching for them.

TRAP **Promoting a second-tier sentence into the first tier.** The commonest version is writing that the review showed the chemical-imbalance explanation is not evidence-based, or that antidepressants do not correct an abnormality. Neither is a result. The result is that a set of literatures did not consistently support one proposition; the two sentences above are what identifiable authors concluded from it, and they were contested in print within months. The fix costs four words: put "the authors concluded that" in front of every sentence in the second tier, and notice that no sentence in the first tier needs it.

THE INFERENCE THAT LOSES MARKS	WHY IT DOES NOT FOLLOW	WHAT WOULD BE NEEDED TO ESTABLISH IT
Antidepressants are ineffective	Efficacy is established in randomised trials against placebo. This review examined evidence about causation and included no efficacy trial at all	The efficacy literature on its own terms: placebo-controlled randomised trials and their meta-analyses
Depression is not biological	One biological account was examined. Nothing about stress-axis regulation, inflammation, neuroplasticity, circuit-level or glutamatergic accounts follows from a null result about serotonin concentration	Evidence against each of those accounts, tested separately and on its own measures
Serotonin has no role in depression	Failure to find a consistent association across the literature collected is not a demonstration that no association exists, and several strands rested on measures that test serotonin function poorly	Adequately powered studies using measures capable of testing serotonin function, reporting null results
Evidence about causation answers a question about treatment	Cause and treatment are separate questions. A drug does not have to reverse the cause of an illness in order to relieve it	Nothing would establish it. This is a category error rather than an unproven claim, which is why it is the costliest of the four

6	2022	1	4
AREAS OF RESEARCH EXAMINED	YEAR OF PUBLICATION	FINDING ASSERTABLE AS THE REVIEW'S OWN	INFERENCES THAT LOSE MARKS

9.4 Why evidence about cause cannot answer a question about treatment

The mechanism by which an effective drug works is not required to be the reverse of the cause of the illness it treats. Analgesia relieves headache and headache is not caused by a shortage of analgesic. The **corrective model**, in which a drug replaces something the body is missing, describes a small minority of effective treatments rather than the general case, which is why the insulin analogy proves less than it appears to. An aetiological null result and a treatment effect are therefore logically independent, and a paper about the first leaves the second where it found it.

Indian guidance puts efficacy on its own footing, which lets this be answered with a source rather than an argument. The 2025 Indian Psychiatric Society depression guideline records that many of the available antidepressants are considered first-line options on the strength of robust randomised placebo-controlled trial evidence. The guideline rests that statement on a network meta-analysis of **21 antidepressant drugs**, which it cites as having confirmed that all these agents are more effective than placebo. Both rest on trials rather than on any theory of causation, and both would read identically if every strand of the umbrella review had come out the other way.

TRAP **Answering an efficacy question with the umbrella review.** A stem asking whether antidepressants work, or whether this controversy changes prescribing, is testing this boundary, and the review is the wrong instrument. The correct move is to say the question is answered by the efficacy literature, name what that literature consists of, then say in one sentence what the umbrella review does bear on: how the drug is explained to the patient rather than whether it is offered. A candidate who cites a review of aetiological evidence as grounds for a prescribing decision has answered a different question well.

9.5 The critiques, and which of them actually bite

Four objections were published, and are not of equal weight. Grade them rather than listing them.

The first is the straw-man objection: that no well-informed academic psychiatrist held a simple chemical-imbalance theory, so a review dismantling it defeats a position the field had already left. Half of that is true, and the Indian guideline says so in its own voice. Its box on the neurobiological mechanisms of antidepressant action opens by recording that antidepressants initially increase the synaptic availability of monoamines such as serotonin and noradrenaline, but that this alone does not explain clinical response. The other half is where the objection fails: the review's target was not only the academic theory but what patients are told, and that account is widely believed and widely used in consultations. A theory abandoned in the seminar room and still spoken in the clinic is a fair thing to examine.

The second is the tool objection, and technically the strongest. An umbrella review is a good instrument for a question a coherent body of systematic reviews was built to answer, and a poor one for assembling several distinct literatures, each measuring a different construct by different methods, into a single verdict on one hypothesis. Critics argued that the included reviews were

heterogeneous in age, quality and purpose, that the quality gradings were contestable, and that some strands, the depletion literature in particular, were read less favourably than the primary studies allow. A multi-author reply published in the same journal, led by Jauhar with several academic psychiatrists, argued along these lines that the serotonin system is implicated in depression and that this method was not capable of showing otherwise. Note what that reply does not do: it contests the inference and the instrument, and it does not restore a simple low-serotonin account.

The third is the logical one, and the cleanest to write. **Absence of evidence is not evidence of absence.** A literature that is underpowered, cross-sectional, heterogeneous in its measures and collected over decades can fail to detect a real effect. Converting the review's failure to find consistent association into a demonstration that no association exists requires a step the design cannot supply. This objection is sound, and it cuts at the authors' broader conclusions rather than at their reported finding, which is why the three tiers above matter.

The fourth concerns consequences rather than method, and should be named without being overweighted. The review was reported widely and in compressed form, and some coverage carried it as a finding that antidepressants do not work. A paper is not answerable for every headline written about it. What belongs in an answer is the clinical consequence, which arrived as patients asking whether they should stop their medication.

> **PITFALL**
>
> Grading a critique by its volume rather than by what it establishes. The loudest objection here is the straw-man one and it is the weakest: it concedes the review's finding and disputes only its novelty. The quietest is the construct objection in the next section, and it is the strongest: it explains why a null result was likely whichever way the biology turns out.

9.6 What the measures could not test, and what was left standing

The strongest move available is not to pick a side. It is to notice that several strands could not have answered the question well even in principle, which weakens the positive literature and the negative conclusion drawn from it in the same movement. Three construct problems recur, and are worth naming in these words.

Concentration is not function: a metabolite concentration in plasma, urine or cerebrospinal fluid is a distant and heavily confounded proxy for serotonin transmission at a synapse in a specific circuit, so a null on the proxy is weak evidence about the thing. Static is not dynamic: the amount of a transmitter present at rest is a different quantity from **release capacity**, the amount a system can release when it is driven, and imaging methods able to probe release with a pharmacological challenge came into use largely after most of the reviewed literature was collected. A positron emission tomography study using an amphetamine challenge, published after the umbrella review, reported reduced serotonin release in a depressed group. It is one small study, it settles nothing, and it is quoted here only to make the structural point: a null covers the measures a literature had, not the measures that came after it. And group is not individual: a

mean difference absent across a diagnostic category assembled from symptoms does not exclude a serotonergic mechanism operating in some subgroup inside it.

Two things the debate never touched should be stated positively; leaving them out makes a script sound as though it has read only the headline. Selective serotonin reuptake inhibitors act on the serotonin system: that is pharmacology, it is not disputed by anyone in this argument, and it is not an aetiological claim. And the neurobiology of depression is broader than any monoamine, with stress-axis, inflammatory, neuroplastic, circuit-level and glutamatergic accounts all under investigation. The Indian guideline's own account is a downstream plasticity account. It puts the typical delay in clinical response at **2 to 4 weeks** and attributes that delay to the time required for gene-level, synaptic and network-level adaptation, stating that antidepressants act by gradually modulating neural circuits rather than by producing immediate relief through neurotransmitter changes alone.

GOOD TO REMEMBER

The most useful thing this controversy leaves behind is not a verdict about serotonin. It is a habit: before repeating any mechanistic explanation of a treatment, ask what would have to be true for it to be right, and whether anyone has checked. That habit transfers to every other explanation offered in a consultation.

9.7 What changes in the consulting room

Nothing about who is offered treatment changes, and everything about how it is explained does. The chemical-imbalance sentence is short, destigmatising and easy to deliver in a crowded outpatient clinic with the family present, which is why it became the default script. It is also a causal claim that the evidence does not support as stated, and it is fragile: a patient who was given it, and who later meets a headline reporting that serotonin was never the problem, has been handed a reason to stop.

The replacement is not "we do not know". Three sentences are defensible and cover what a patient is actually asking. This medicine acts on the serotonin system, which is a statement about the drug rather than about the illness. Why that produces improvement is not fully established, and the current account is about gradual adaptation in brain circuits rather than the topping up of a missing chemical. And the reason to expect benefit is that the medicine has been compared against placebo in randomised trials, which is a different kind of evidence from any theory of cause. Add the expected time to effect and the common adverse effects, and it still fits inside a consultation.

IN INDIA

The 2025 Indian Psychiatric Society depression guideline directs that before prescribing, clinicians provide psychoeducation covering expected benefits, possible risks, anticipated time to effect, and the addressing of myths or misconceptions about antidepressant use, and in its section on suicidal behaviour it directs that the patient and, with their consent, family members be educated about the bio-psycho-social model. Neither instruction endorses a chemical-imbalance script and the second is explicitly broader than any neurochemical account. Note also what is not there. A whitespace-insensitive text search of the four 2025 Indian Psychiatric Society guidelines held for these notes returns no occurrence of "chemical imbalance", "serotonin hypothesis", "serotonin theory" or "monoamine hypothesis" in any of them, while "serotonin" itself occurs 17 times in the depression guideline alone. That is a search of these copies rather than a statement about Indian practice: no Indian guideline position on this controversy, and no Indian survey of what patients are told about the mechanism, is held by this account, and neither should be invented in an answer.

9.8 Answering the question, whatever form it takes

Reduced to five moves, the topic answers as an essay, a short note or a viva.

- Name the design and its inheritance: an umbrella review of systematic reviews, quality-rated, its reach limited by what it collected.
- Give the six strands with a finding attached to each, and mark the sixth as pointing the other way.
- State the reported finding flatly, then attribute every broader conclusion to the authors by name.
- Grade the critiques instead of listing them, and say which would survive a result in the opposite direction.
- Close on the boundary: this is evidence about cause, the efficacy question is answered elsewhere, and what changes is the explanation given to the patient.

MNEMONIC

Design, Data, Authors, Critics, Boundary

Five steps in one direction, and the order is the argument. Design constrains what the data can show; the data are separable from what the authors made of them; the critics attack the middle three and not the first tier; and the boundary is where the marks are, because it is the step everybody skips.

The review reported that a literature does not support a claim; everything stronger than that belongs to somebody, and an answer that says whose is the answer that scores.

10 · Psychedelic-assisted therapy (psilocybin, MDMA) and regulatory status

The examinable object here is not a molecule, it is a package: psilocybin-assisted therapy and MDMA-assisted therapy. Both are a drug given inside a structured psychological intervention with four components. Half the marks sit on reading the evidence honestly: expectancy, functional unblinding, therapist effects, and who was screened out. The other half sit on regulatory status, where several arrangements that read like approval are not one, and only a marketing authorisation licenses a prescription. The pharmacology of these substances as agents of intoxication and misuse belongs to the Substance use disorders notes; this item covers the therapeutic trial and the regulatory status of both agents in three named jurisdictions.

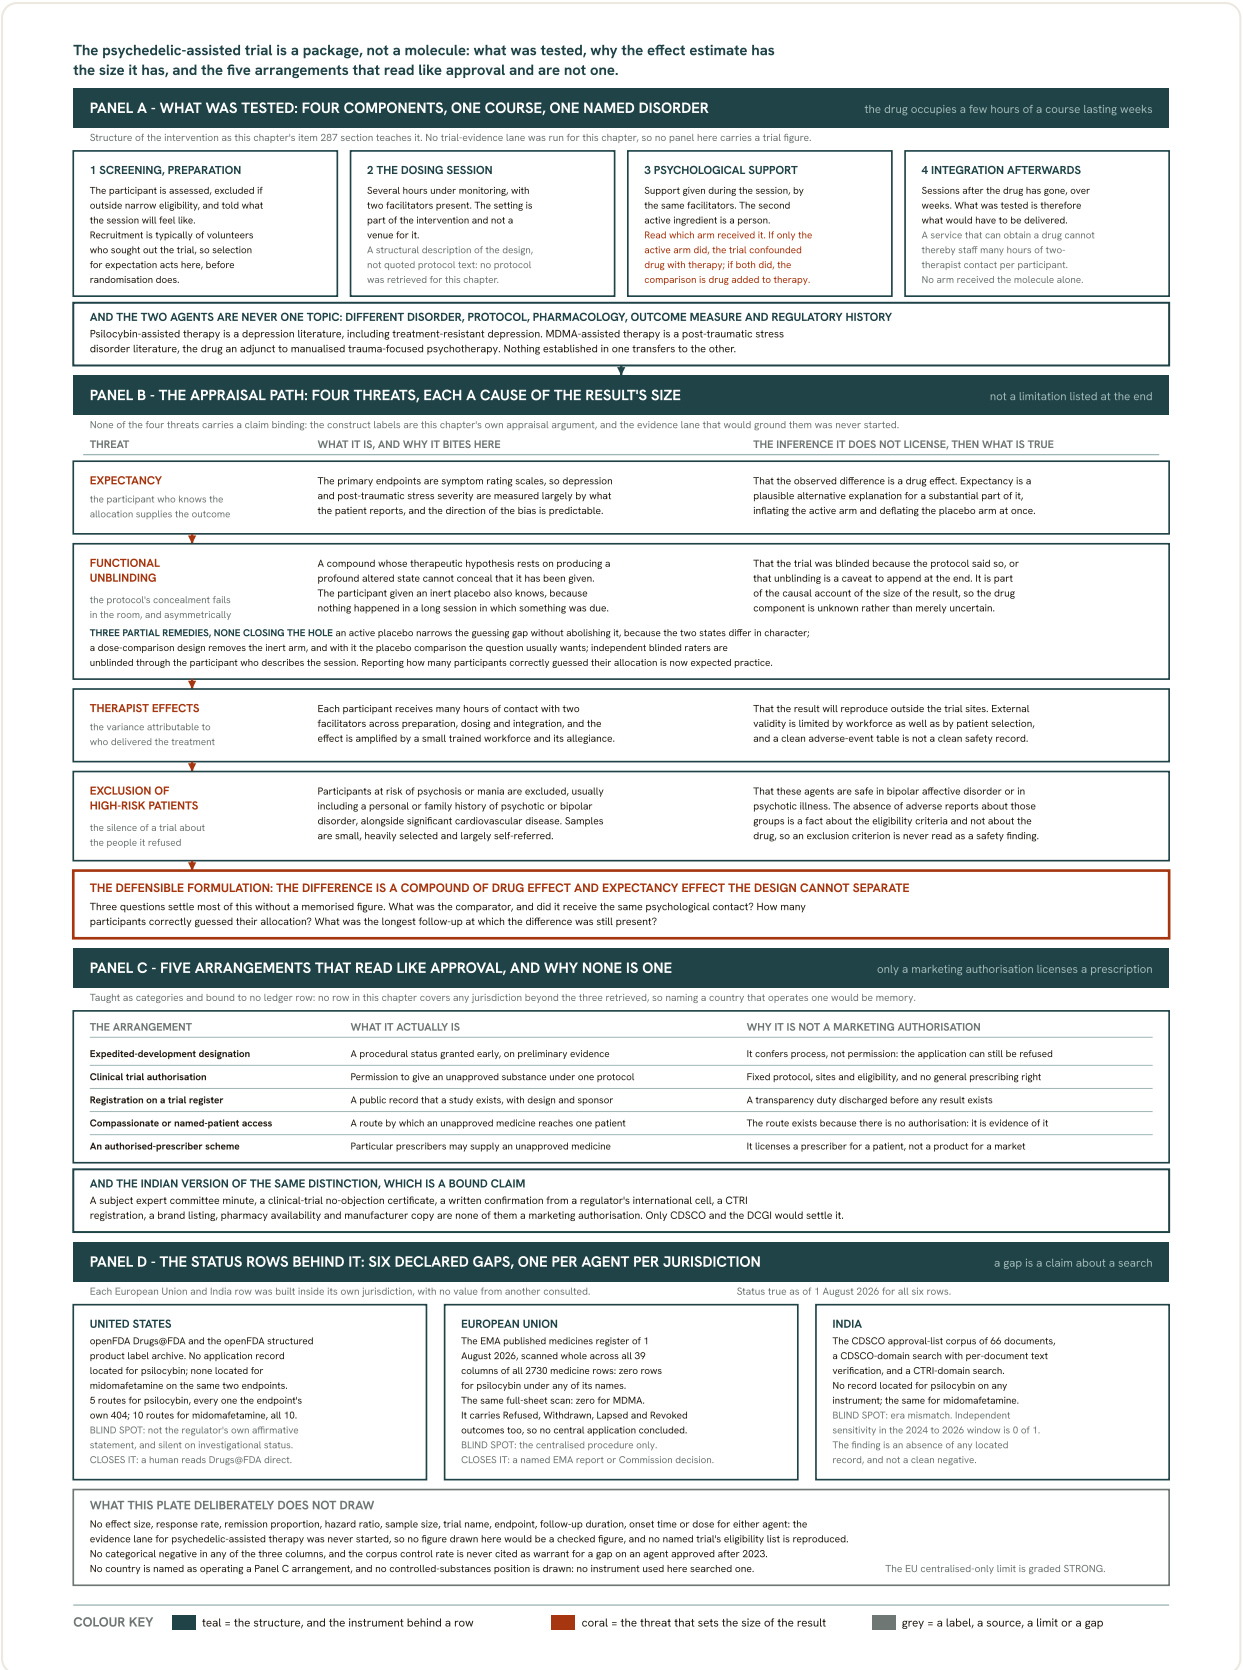


Fig 38.6 Appraising a psychedelic-assisted trial: the package tested, the four threats that set the result's size, and the arrangements that are not a marketing authorisation. The unit of evidence is the package, never the molecule. The four threats are causes of the effect estimate, not caveats after it, so an exclusion criterion is never a safety finding. No effect size, sample size, trial name or follow-up figure is drawn: that lane was never started. Six declared gaps, each naming its instrument: openFDA Drugs@FDA and the label archive, the EMA register scanned whole, and the CDSCO approval-list corpus with CDSCO-domain and CTRI-domain searches, true as of 1 August 2026. (Day 38 status ledger, psilocybin and midomafetamine; the appraisal structure is this chapter's own.)

10.1 What is actually being tested, and why the answer is never "the drug"

Almost every error here is an error of abbreviation. The **intervention package** has four components: screening and preparation, in which the participant is assessed, excluded if outside narrow eligibility, and told what the session will feel like; the dosing session, several hours with monitoring and two facilitators; the psychological support given during it; and integration afterwards. The drug occupies a few hours of a course lasting weeks.

Two consequences follow. The first is attributional: no arm received the molecule alone, so the result cannot be attributed to it, and the intervention's name is psilocybin-assisted therapy, not psilocybin. The second is transfer: what was tested is what would have to be delivered, and a service that can obtain a drug cannot thereby staff many hours of two-therapist contact per participant.

The two agents should never be taught as one topic. The psilocybin literature is a depression literature, including treatment-resistant depression. The MDMA literature is a post-traumatic stress disorder literature, the drug an adjunct to manualised trauma-focused psychotherapy. Different disorder, protocol, pharmacology, outcome measure and regulatory history. Nothing established in one transfers to the other, and a stem that mixes them is testing exactly that.

RULE **Name the package, not the molecule, and never merge the two agents.** The unit of evidence is drug plus preparation plus supported session plus integration, in one named disorder, against one named comparator, over one stated follow-up period. Any sentence that drops a component, or carries a psilocybin depression finding to MDMA in post-traumatic stress disorder, has changed the thing being claimed.

10.2 Two mechanisms, and the reason the trial cannot be blinded

Psilocybin is a prodrug, dephosphorylated to psilocin, an agonist at the **5-HT2A** receptor; the acute state is a marked alteration of perception, affect and self-experience lasting several hours. MDMA is a substituted amphetamine acting principally as a monoamine releaser, with prominent serotonin release, producing emotional openness, reduced defensiveness and sympathomimetic arousal. Its rationale is not that it treats the disorder but that it makes trauma-focused psychotherapy tolerable for people who cannot otherwise engage with the memory.

Read those descriptions as design constraints. A compound whose therapeutic hypothesis rests on producing a profound altered state cannot conceal from the participant that it has been given, so the trial is blinded on paper and unblinded in the room, and every downstream problem descends from it.

These are not agents given daily to steady state. The exposure is one or two supervised sessions, and the benefit is claimed to outlast the drug by weeks or months. That puts a heavy burden on follow-up duration: a result at four weeks says nothing about six months.

10.3 Expectancy and functional unblinding, and what they do to a subjective endpoint

Blinding is not a formality here; it is the load-bearing wall, and it is cracked. **Functional unblinding** means the allocation concealment specified in the protocol fails in practice because the treatment announces itself. It fails almost completely, and asymmetrically: the participant

given an inert placebo also knows, because nothing happened during a long session in which something was expected.

The primary endpoints are symptom rating scales: depression and post-traumatic stress severity are measured largely by what the patient reports. So the participant who knows their allocation also supplies the outcome measurement. **Expectancy** is not a small additive nuisance there; it is a plausible alternative explanation for a substantial part of the observed difference, and its direction is predictable, inflating the active arm and deflating the placebo arm at once.

Three partial remedies exist and none closes the hole. An **active placebo**, producing noticeable but non-psychedelic effects, narrows the guessing gap without abolishing it, because the character of the two states differs. A dose-comparison design, every arm receiving the drug at a different dose, removes the inert arm and the sharpest expectancy contrast, but also the placebo comparison the question usually wants. Independent blinded raters improve on self-report, but a rater interviewing a participant describing their session is unblinded through the participant. Reporting how many participants correctly guessed allocation is now expected practice.

TRAP **Treating unblinding as a limitation to be listed at the end.** It is not a caveat appended to the result; it is part of the causal account of the result's size. The defensible formulation: the between-group difference is a compound of drug effect and expectancy effect the design cannot separate, so the size of the drug component is unknown rather than merely uncertain.

10.4 The therapist is inside the intervention, not outside it

The second active ingredient is a person: each participant receives many hours of contact with two facilitators across preparation, dosing and integration, emotionally charged throughout. **Therapist effects**, the variance in outcome attributable to who delivered the treatment rather than what was delivered, are well recognised in psychotherapy research and amplified here by the small number of trained facilitators, their allegiance to the model, and the intensity of the dyad during an altered state.

Three examinable consequences follow. First, attribution again: if only the active arm receives many hours of skilled psychological contact the trial has confounded drug with therapy; if the placebo arm receives the same contact the comparison is drug added to therapy, not therapy itself. Read which the trial did. Second, replication: delivery depending heavily on individual facilitators is hard to reproduce outside the trial sites, so external validity is limited by workforce as well as patient selection. Third, adverse-event ascertainment: silence collected by an instrument that cannot see a harm is not evidence that the harm did not occur.

PITFALL

Reading a clean adverse-event table as a clean safety record. The suggestibility and dependence on the dyad the model relies on therapeutically are the same features that create the risk of boundary violation, so mechanism and characteristic harm share a root. A safety discussion that lists nausea, headache, transient blood pressure rise and anxiety has described the physiology and omitted the setting.

10.5 Who was enrolled, who was screened out, and what the result therefore covers

Samples here are small and heavily selected. Recruitment is typically of volunteers who sought out the trial, which selects for expectation before randomisation acts. Screening is intensive. And, as a standing feature of the field, participants at risk of psychosis or mania are excluded, usually including a personal or family history of psychotic or bipolar disorder, alongside significant cardiovascular disease.

These trials are silent about the people they excluded. There is therefore no evidence from this literature that psychedelic-assisted therapy is safe in bipolar affective disorder or in psychotic illness, and the absence of adverse reports about those groups is a fact about the eligibility criteria, not about the drug. A candidate who writes that these agents "appear safe, with no cases of induced psychosis reported" has read an exclusion criterion as a safety finding. The same logic bounds every claim: the result describes the population enrolled, the protocol used, the comparator used and the follow-up point reported, and nothing else.

This account prints no effect size, response rate or remission proportion for either agent, and the omission is deliberate. The trial-evidence retrieval for this chapter was not completed, so no figure here would be a checked figure, and a remembered number in this field is characteristically right about the wrong thing: right for a different trial, comparator, endpoint or week. Attach a number only from the paper in front of you, with its population, comparator and time point.

GOOD TO REMEMBER

Three questions settle most viva exchanges here, none needing a memorised figure. What was the comparator, and did it receive the same psychological contact? How many participants correctly guessed their allocation? What was the longest follow-up at which the difference was still present?

10.6 The status ledger for psilocybin and MDMA, jurisdiction by jurisdiction

Regulatory status is recorded on four fields, and a status sentence exists only where a named regulator's record was retrieved: **jurisdiction**, the **exact indication**, the **authoritative source**, and the date the status was true. For both agents, in all three jurisdictions examined, the retrieval found no record. That is six declared gaps, and the correct sentence is not "these agents are not approved" but what was searched, on whose authority a claim would rest, and when.

AGENT AND JURISDICTION	WHAT WAS SEARCHED	WHAT WAS FOUND	STATUS TRUE AS OF
Psilocybin, United States	openFDA Drugs@FDA and the openFDA structured product label archive, 5 query routes, each positive-controlled in the same run	No application record for generic name psilocybin; every route returned the endpoint's own HTTP 404, zero transport failures	1 August 2026
Midomafetamine (MDMA), United States	The same two openFDA endpoints, **10 query routes**, probing midomafetamine, midomafetamine hydrochloride, the full chemical name and the brand-name field	No application record for generic name midomafetamine; all 10 returned HTTP 404, zero transport failures	1 August 2026
Psilocybin, European Union	The EMA published medicines register generated 1 August 2026, scanned whole across all 39 columns of all 2730 medicine rows	Zero rows under any of its names; the register carries Authorised, Refused, Application withdrawn, Withdrawn, Lapsed, Expired, Revoked, Suspended and Opinion outcomes, so this is not merely "not authorised"	1 August 2026
Midomafetamine (MDMA), European Union	The same full-sheet scan of the same register	Zero rows under any of its names; no EU centralised procedure of any kind is on record for this active substance	1 August 2026
Psilocybin, India	The CDSCO published approval-list corpus of 66 documents, a CDSCO-domain search with per-document text verification, and a CTRI-domain search with a positive control	Absent from all 66 lists; the CDSCO-domain search returned zero results at exit status 0, a genuine empty result set rather than a failed query; no indexed CTRI trial record	1 August 2026
Midomafetamine (MDMA), India	The same three Indian instruments	The same finding as psilocybin: no CDSCO record of any kind located on any instrument, and no indexed CTRI trial record	1 August 2026

Two features of that table matter. First, each row was built inside one jurisdiction only: absence in the European Union column was established from EMA sources alone and absence in the India column from Indian sources alone, with no value from another jurisdiction consulted, because a

row may be authored only in the jurisdiction it was grounded in, a declared gap is authorable as a declared gap, and an inferred figure is not authorable at all. Second, the instruments were proved working before their silence was believed: four of four hosts round-tripped a known-present record before any miss was read as a fact, with certificate verification never disabled. An earlier run reported nothing anywhere for any agent; the cause was a broken certificate store, not a fact about a drug.

2730	66	10	0
EU REGISTER ROWS SCANNED	CDSCO APPROVAL DOCUMENTS READ	OPENFDA ROUTES, MIDOMAFETAMINE	RECORDS LOCATED, ALL JURISDICTIONS

Each instrument has a measured blind spot. The American search covers openFDA Drugs@FDA and the label archive only; it is not the regulator's own affirmative statement that no application exists, and it is silent on investigational and breakthrough designations, which appear in neither instrument. That silence is the trap: the same null comes back whatever stage a programme has reached, so it is no evidence that nothing is under way. What either instrument indexes, and what it leaves out, is a property of the instrument, and this chapter did not search it: what was measured is the scope sentence above and nothing wider. A claim about the coverage of either endpoint would rest on the regulator's own description of it, retrieved and dated, and none was retrieved here. Closing the row: a human must read Drugs@FDA directly and, if there is no application, say so on the regulator's authority rather than on this probe's silence.

The European search covers the centralised procedure. A medicine authorised nationally through the decentralised or mutual-recognition route would not appear in the register at all, so absence here means no EU centralised procedure of any kind on record, which is narrower and weaker than "not available anywhere in the European Union". That limit was graded: STRONG for both, because for a small molecule outside the mandatory centralised scope a national authorisation is genuinely possible and wholly unmeasured here. A second weakness: the confirming second instrument was applied to 8 of 8 grounded rows and 0 of 8 gap rows, so it fired only where the first instrument had already found something, and no European gap in this chapter has independent corroboration. A claim here would have to rest on an EMA European public assessment report or a European Commission implementing decision, named as such; the phrase "approved internationally" may never be substituted for a named regulator.

The Indian instrument is the weakest of the three, and its weakness is measured, which is why the India cells are gaps and not negatives. Of 20 drugs taken as approved and marketed in India, 17 appear in the corpus and 3 do not: agomelatine, molnupiravir and trospium. More decisive is the era mismatch. Every one of the 17 present controls appears in a document dated 2023 or earlier, independent sensitivity in the 2024 to 2026 window is 0 of 1, and no sentence may cite the 17 of 20 figure as warrant for a gap on any agent approved after 2023. The corpus also contradicts itself: CDSCO's own subject expert committee minutes of 16 April 2025 state that another agent in this chapter, whose status row is another item's, was already approved in the country, yet that agent appears in none of the 66 lists. Hence the corrected wording on both India rows: "clean negative" overstates what an instrument with a measured 15 per cent miss rate, blind to three of

its own positive controls, can deliver, and the finding is an absence of any located record rather than a clean negative.

TRAP **Writing "psilocybin is not approved in India".** That claims knowledge of the regulator's whole record, and this instrument cannot support it. What it supports is narrower and checkable: no CDSCO approval record for psilocybin was located in the 66 published approval lists, in a CDSCO-domain search that returned a genuine empty result set, or in a CTRI-domain search, as of 1 August 2026. The general form: name the register, name the date, name what would close the row, and let the absence be an absence.

The falsification pass that tested this chapter's status rows against their primary documents scored all six not applicable, on the stated ground that a declared gap carries no indication text, so an instrument that compares text against a primary document has nothing to test: a gap is not a claim about a product, it is a claim about a search. They have been checked for what they say about the search, never independently confirmed as absences.

10.7 Five things that are not a marketing approval

Each of the five below is routinely written up, in press and review articles, in language indistinguishable from approval. A marketing authorisation is a decision by a named regulator that a specific product may be placed on the market for a stated indication.

MECHANISM	WHAT IT ACTUALLY IS	WHY IT IS NOT APPROVAL
Expedited-development designation	A procedural status granted early, giving a sponsor closer regulatory interaction and faster review of a programme for a serious condition	Granted on preliminary evidence, it confers process, not permission: the application can still be refused or withdrawn, with no product reaching the market
Clinical trial authorisation or permission	Permission to give an unapproved substance to defined participants under a defined protocol at defined sites	It authorises research, and its conditions are the opposite of a marketing licence: fixed protocol, sites and eligibility, no general prescribing right
Registration on a trial register	A public record that a study exists, with its design, sponsor and eligibility criteria	A transparency requirement discharged before results exist. It says a study was declared, not that it worked or that anything was licensed
Controlled, compassionate or named-patient access	A route by which an unapproved medicine reaches a specific patient with no alternative, usually with individual authorisation and reporting conditions	The route exists precisely because there is no marketing authorisation. It is evidence of the absence it works around
An authorised-prescriber or special-access scheme	An arrangement permitting particular approved prescribers to supply an unapproved medicine to particular patients	It licenses a prescriber for a patient, not a product for a market. The product remains unapproved, and the scheme's paperwork usually says so

Two distinctions belong beside it. First, controlled-substances law, on which this chapter searched nothing. Every instrument behind these notes is a drug-approval or drug-trial instrument, and not one is a statute, a scheduling notice or a controlled-substances register, so nothing retrieved here establishes what any schedule permits or forbids, what altering one would do, or how a schedule stands beside a marketing authorisation in any jurisdiction. A claim about any of that would rest on the statute itself and on the body that administers it, named, quoted and dated. Until one is read, treat the schedule and the licence as two questions put to two records, and answer only the one these notes hold a record for. Second, chapter hygiene: a subject expert committee minute, a clinical-trial no-objection certificate, a written confirmation from a regulator's international cell, a CTRI registration, a brand listing, pharmacy availability and manufacturer copy are none of them a marketing authorisation, and the only bodies whose record would settle the Indian question are CDSCO and the Drugs Controller General of India.

This section names no country as operating any of the five mechanisms, deliberately: no ledger row covers any jurisdiction beyond the three retrieved, so attaching a named scheme to a named country would be writing from memory.

IN INDIA

The position to hold, and it is narrow. As of **1 August 2026**, no CDSCO or DCGI approval record for psilocybin was located in the 66 published approval lists, in a CDSCO-domain search, or in CTRI, and the finding for MDMA is identical. Say that rather than "not approved", and say what would change it: an affirmative CDSCO or DCGI new-drug permission, a current CDSCO-approved label, or an entry in an official approval list. A CTRI registration, if one appears, establishes that a study exists and does not close the approval question. No instrument used for this chapter looked at the controlled-substances position for either agent in India, so nothing here establishes what that position is or what it requires, and the schedule number stays out of your answer until the statute and the body that administers it have been read. And when a patient or family arrives having read about these treatments abroad, name the gap rather than filling it: no located Indian marketing authorisation for either agent, a multi-session therapy package rather than a prescription, and trials that excluded several groups who reach Indian outpatient departments in large numbers.

10.8 What a complete answer contains

Four moves answer this item, in any form.

- Define the intervention as a package, name its four components, and attribute the result to the package, not to the molecule.
- Give functional unblinding and expectancy as the reason the effect estimate has the size it has, not as a limitation appended afterwards, and say what a design can and cannot do.
- Bound the evidence by the population enrolled: small, selected, self-referred, screened to exclude those at risk of psychosis or mania, so no safety claim is available for bipolar or psychotic illness.

- State regulatory status on the four fields, register named, date attached, gap declared, and separate approval from designation, trial permission, registration, access arrangement and prescriber scheme.

MNEMONIC

Jurisdiction, Indication, Source, Date

The four fields of a status sentence, in writing order. A claim missing any one is not weaker, it is different: without jurisdiction a claim about the world, without indication a claim about a molecule, without a source a claim about the writer's memory, without a date a claim about now, forever.

RULE **Absence anywhere in this section is a statement about a search, never proof of absence.** A queried register's null is evidence and worth stating; it is not proof, and the difference has been demonstrated here. One agent in these notes, brexanolone, was recorded as having no United States record until multi-route probing retrieved application NDA211371, ZULRESSO, approved 17 June 2019 and now of Discontinued marketing status: the record carries no openfda block at all, so three of five query routes returned NOT_FOUND while a fourth returned it. Strengthen a null by widening the routes, not by repeating one: an independent recheck of the psilocybin row on a separate run repeated the probe across 9 routes, including free-text and wildcard searches on both endpoints, returned zero on every one, and had its own positive controls return hits in the same run.

The package is what was tested, unblinding is why the number is uncertain, exclusion is what the number does not cover, and a designation, a trial permission and an access scheme are not a licence.

11 · Artificial intelligence and machine learning in psychiatry

A model does not diagnose anybody; it returns a number, and the whole of this topic is what that number is allowed to mean. The marks sit not on architectures but on four questions asked of any reported figure: what was the model predicting, in whom, against what, and had the data ever been seen before. Candidates lose them in three predictable places: an accuracy figure quoted without its population, a good discrimination statistic read as a good decision, and a performance claim treated as a regulatory permission. The statistics themselves, sensitivity, specificity, predictive values and the appraisal of a diagnostic accuracy study, belong to the research methodology and biostatistics notes; what follows is only what changes when the index test is a fitted model rather than an assay.

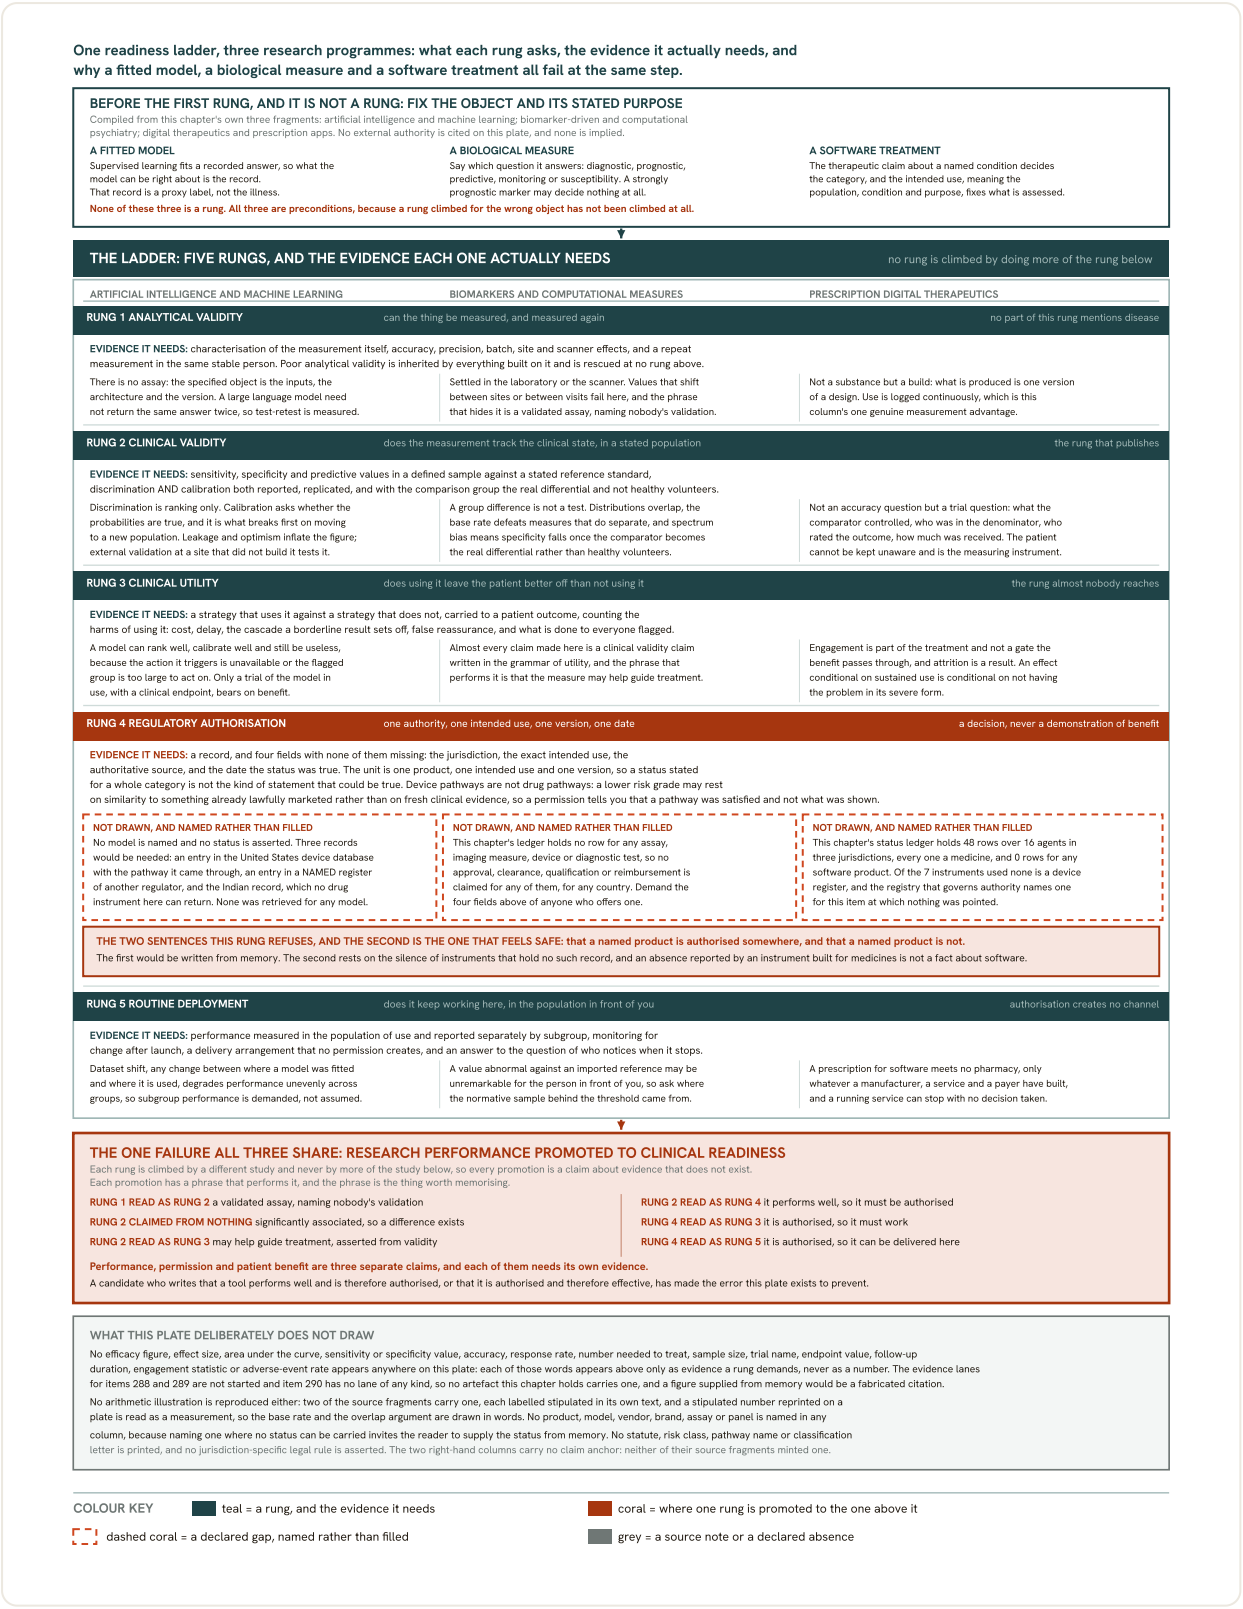


Fig 38.7 One readiness ladder for artificial intelligence, biomarkers and prescription digital therapeutics: analytical validity, clinical validity, clinical utility, regulatory authorisation, routine deployment. Above the first rung the object and its purpose are fixed: supervised learning fits a recorded answer, and that record is a proxy label rather than the illness. Each rung is climbed by a different study, never by more of the study below, so discrimination and calibration settle rung 2 and neither touches rung 3, and external validation at a site that did not build the model answers what leakage and optimism flatter. Rung 4 is drawn EMPTY and named rather than filled: no product is named and no status asserted. Rung 5 is where dataset shift bites. Source: this chapter's own three fragments, items 288, 289 and 290; the evidence lanes for 288 and 289 read as not started and item 290 has no lane at all.

11.1 What the machine is actually predicting

Almost every clinical model here is **supervised learning**: the algorithm is shown examples in which the answer is already recorded, and fits a function from the inputs to that recorded answer. Everything it can ever be right about is therefore the record, not the illness, the first thing to establish about any reported result and the one most often skipped.

The recorded answer is a **proxy label**. In psychiatry the proxies available at scale are administrative or documentary: a diagnostic code entered at discharge, a prescription of a particular class, a rating-scale score above a cut-point, a chart entry flagging risk, a readmission within a stated window. None of these is the construct: a code records what a clinician wrote down in a system that pays for codes. A model trained on codes learns coding practice, which varies by clinician, by unit and by country, and reproduces it faithfully.

The construct validity of the label is a ceiling on everything downstream: no split, no architecture and no sample size repairs a target that does not mean what the paper says it means. And the label determines what the number transfers to, so a model that predicts a recorded diagnosis of depression predicts documentation, and a claim that it predicts depression has quietly changed the target between the methods section and the abstract.

11.2 The six things a performance number carries, or it is not a finding

A performance figure with nothing attached is uninterpretable, not merely imprecise. Six items travel with it, each naming a distinct way a good-looking number goes wrong.

WHAT MUST BE STATED	WHY IT IS NOT OPTIONAL	WHAT A MISSING ENTRY HIDES
The model	Inputs, architecture and version define the thing being reported	A retrained or updated model is a different model with a different number
The target	The exact label, how it was defined and by whom	The proxy problem above, and a target that drifts between sections
The population and setting	Who was included, recruited how, where and when	Case mix and prevalence, which move predictive values without moving the model
The split method	How the evaluated data were separated from the fitted data	Leakage and optimism, the commonest cause of an inflated figure
The comparator	What the model is being judged against	A model beating nothing, or beating a deliberately weakened baseline
External validation status	Whether the number comes from data the model never saw, from a site that did not build it	The single largest source of the gap between published and delivered performance

The comparator row is the one candidates forget and an examiner can mark hardest. A model that predicts an outcome better than chance has been compared with nothing. The comparisons that matter are the ones a service would otherwise use: the existing clinical judgement, an established

structured instrument, or a simple regression on a handful of routinely available variables. That last comparator is unglamorous and frequently wins, which is why a paper omitting it has omitted the informative arm.

RULE **A performance figure belongs to one model, one target, one population, one split, one comparator and one validation status, and it does not travel without all six.** Move any one of them and the number is no longer about the same thing. This is the whole discipline of the topic, and why a sentence of the form "artificial intelligence can detect depression with high accuracy" is not a weak claim but an empty one: it names none of the six.

11.3 Discrimination is not calibration, and neither one is usefulness

Discrimination is the model's ability to rank: given one person who has the outcome and one who does not, how often does it give the first the higher score. That is what the area under the receiver operating characteristic curve measures, and its floor is worth fixing in memory: a model at **0.5** ranks no better than a coin. Discrimination is a property of ordering only, and says nothing about whether the predicted probabilities are true.

Calibration is that second question: among the people the model gave a twenty percent probability, did about twenty percent have the outcome. A model can rank well and be badly calibrated, systematically over-predicting risk in everybody. The distinction has a direct clinical consequence: a threshold is set on the probability, not on the ranking, so a poorly calibrated model puts the operating threshold in the wrong place even when its ranking statistic looks strong. Calibration is also the property most damaged by moving a model to a new population, and the least often reported.

Neither statistic answers the question a service actually has: whether acting on the output leaves patients better off than not acting. That is a question about consequences at a chosen threshold: what is done to the people flagged, what it costs, and what is missed in the people not flagged. Decision-analytic measures of net benefit address it, and no ranking statistic substitutes. A model can discriminate well, calibrate well and still be useless, because the action it triggers is unavailable, or the group flagged is too large to act on.

11.4 The base rate, and why an accurate model flags mostly false alarms

This is the most productive arithmetic in the topic, and where risk-prediction claims in psychiatry come apart. Sensitivity and specificity are properties of the model. Predictive values are not: they depend on how common the outcome is in the people the model is applied to, and psychiatry's most-wanted outcomes are rare. What follows is a worked example with numbers chosen for arithmetic, not measured in any study, and worth reproducing in an answer as a worked example.

Take an outcome occurring in **1 in 100** of those screened, and a model with **80 percent** sensitivity and 80 percent specificity, which would be a respectable psychiatric model. In 10,000 people, 100 have the outcome and the model finds 80 of them. Of the 9,900 who do not, it wrongly flags 1,980. So 2,060 people are flagged, 80 of them with the outcome.

1%	80%	3.9%	99.7%
ASSUMED BASE RATE, NOT MEASURED	ASSUMED SENSITIVITY AND SPECIFICITY	RESULTING POSITIVE PREDICTIVE VALUE	RESULTING NEGATIVE PREDICTIVE VALUE

Read both ends of that strip. About 26 people are flagged for every one who has the outcome, so whatever the flag triggers is delivered, and paid for, overwhelmingly in people who will not have the event. Meanwhile the negative predictive value looks magnificent and is uninformative: with an outcome this rare, predicting no event in everybody already achieves 99 percent, so the model must be judged against that baseline and not against zero. Both figures move whenever the model is applied to a population with a different base rate, which is why a screening claim from a specialist inpatient sample does not survive the move to a general clinic.

PITFALL

Quoting accuracy as the headline figure when the outcome is rare. Accuracy is the proportion of all predictions that were correct: with an outcome occurring in one percent of people, a model predicting no event for everybody is 99 percent accurate and clinically worthless. An answer repeating an accuracy figure on an imbalanced outcome without noticing the imbalance has walked into the oldest trap in the field.

11.5 How the number gets inflated: leakage, optimism and the wrong split

Published performance is usually higher than delivered performance, and the reasons are structural rather than dishonest. **Data leakage** is information about the answer reaching the model through a route that will not exist in use. Its forms are separate examinable errors: preprocessing such as scaling or feature selection performed on the whole dataset before the split, so the held-out data influenced the fitting; a feature that is a consequence of the outcome rather than a predictor of it, such as a treatment started because the event had already occurred; the same patient contributing records to both sides of the split; and a time window that lets the model see data recorded after the moment the prediction is supposed to be made.

Optimism is the second mechanism and it needs no error at all. When many models, many feature sets and many thresholds are tried and the best result is reported, that result is partly a selection from noise, and it shrinks on new data. The same happens when the held-out set is consulted repeatedly during development, at which point it has joined the training process. The problem is worst where candidate features greatly outnumber patients, the ordinary condition of neuroimaging and genomic work in psychiatry.

Both are properties of how the evaluation data were separated, which makes the split method the most informative line in a methods section. The ladder runs weakest to strongest.

SPLIT	WHAT IT HOLDS OUT	WHAT IT STILL CANNOT RULE OUT
Random split of records	Individual rows	The same patient on both sides, and every form of leakage
Patient-level split	Whole patients	Drift over time, and anything peculiar to that one site
Temporal split, trained on earlier data and tested on later	The future of one site	Everything specific to that site: case mix, coding habits, systems
External validation at a site that did not build the model	Site, population and practice	That the model, once acting on care, still performs as measured
Prospective evaluation with the output withheld from clinicians	Retrospective advantage	Whether acting on it changes any patient outcome

Even the bottom rung falls short. A model can be externally validated, prospectively confirmed and well calibrated and have no demonstrated effect on any patient outcome, because performance and benefit are different endpoints and only a trial of the model in use measures the second. Reporting and appraisal frameworks exist and are worth naming: TRIPOD plus AI for the reporting of prediction model studies, PROBAST plus AI for risk-of-bias appraisal, SPIRIT plus AI and CONSORT plus AI for protocols and reports of trials of these models, and DECIDE plus AI for early live clinical evaluation. The mark goes to knowing what each covers, since a paper satisfying the reporting standard is completely reported and not thereby correct.

TRAP **Reading an internal cross-validation figure as though it were an external result.** Cross-validation resamples inside one dataset, so it estimates how the fitting procedure behaves there and nothing else. It does not test a different site, a different case mix, a different coding culture or a different era, and where preprocessing or feature selection was done before the folds were drawn it does not even estimate that honestly. Examinably: internal validation bounds optimism within one sample, and only evaluation on data from an institution that had no part in building the model bears on whether it transfers.

11.6 Large language models: a different tool with a different failure surface

A **large language model** is not a clinical classifier and should not be appraised as one. It is trained to produce plausible continuations of text, so its outputs are fluent by construction and true only incidentally, and the failure modes follow from that rather than from any defect in a particular product. It produces confident, well-formed statements that are false, including references that do not exist. It tends to agree with the framing it is given, so a leading clinical question returns a supporting answer. It emits no calibrated probability alongside its answer, removing exactly the quantity the previous sections were about. And it need not return the same answer twice, so test-retest reliability cannot be assumed and has to be measured.

Two further points are examinable because counter-intuitive. Performance on written examination questions is a poor guide to clinical performance, partly because a question with five options and one defensible answer is not the task of the clinic, and partly because material that has circulated on the internet may sit in the training data, so a benchmark score can reflect

exposure rather than reasoning. And a general-purpose model given clinical text is not thereby a medical device or a validated instrument; what it is depends on the claim made for it and how it is deployed.

GOOD TO REMEMBER

The nearest-term uses in psychiatry are administrative rather than diagnostic: drafting documentation from a recorded consultation, summarising a long record, coding, translation support. These carry real benefit through time returned to the clinician, and a distinct risk profile, since a fabricated or subtly altered detail in a summary enters the record and is read as fact downstream. The safeguard is structural rather than technological. Output is a draft; the clinician who signs it owns every word in it, and the review has to be real, because a summary that is right nine times out of ten trains the reviewer to stop reading it.

11.7 Regulatory status: what this section refuses to say, and why

Nothing in this fragment states that any named model, product or system is approved, cleared, licensed or authorised anywhere. That is a deliberate refusal and the correct examination position, because the approval status of software is per product, per intended use and per version, and there is no class-wide approval of artificial intelligence in medicine of which any individual tool is an instance. The status ledger built for this chapter covers medicines only, and a drug ledger licenses no statement about a device.

Three records would have to be held before any regulatory sentence could be written about a named model, and naming them demonstrates the point better than reciting a status. The first is an entry in the United States device database for that specific product with the pathway it came through, since clearance against a predicate device, authorisation as a novel low-risk device and full premarket approval are different standards with different evidence behind them. The second is an entry in a named register of another regulator, identified rather than assumed from the first, because a marking obtained in one jurisdiction is not a permission in another. The third is the Indian record, and it is where an answer should be most careful, because the instruments this chapter searched are drug instruments: nothing retrieved here establishes which Indian pathway governs a software product, which body administers it, or what class and licence any given product carries. Those are matters of record, and this chapter holds no such record and did not look for one. No such record was retrieved for this chapter for any model, so no product is named here and no status is asserted for any.

Two errors of reasoning persist even when the records exist. A regulatory permission is evidence that a standard was met for a stated intended use; it is not evidence that patients do better, and the evidence supporting a clearance is often accuracy against a reference standard rather than any outcome. And because the object approved is a version, a model that continues to learn, or is retrained, raises the question whether the thing in use is still the thing assessed, which is why change-control for adaptive software is a live regulatory problem rather than a settled one.

TRAP **Sliding from a performance claim to a permission, or from a permission to a benefit.** These are three separate objects and each transition is a separate error. A published accuracy figure is a measurement in a sample; a regulatory clearance is a decision about one product for one stated use; a demonstrated patient benefit needs a prospective study of the model in use with a clinical endpoint. A candidate who writes that a tool is approved and therefore effective, or that it performs well and is therefore approved, has made the error the whole topic exists to prevent.

11.8 Bias, the Indian question, and what a complete answer contains

A model cannot be called fair, and the reason is that fairness is not one quantity. Several formal definitions of fairness exist and conflict mathematically, so a model satisfying one will generally violate another, and a claim of fairness without a stated definition and a measurement is not a finding. The mechanisms producing unequal performance are, however, straightforward. Groups under-represented in the training data are fitted less well. The label itself is socially patterned, since who receives a recorded diagnosis, who is flagged as a risk and who is admitted all depend on who is doing the recording. And **dataset shift**, meaning any change between the setting where the model was fitted and the one where it is used, degrades performance in ways that fall unevenly across groups. The remedy in an answer is not an assurance but a demand: performance reported separately by subgroup, in the population of use.

IN INDIA

The motivation is real and the strongest argument in the topic: the specialist workforce is far smaller than the need, and any tool that extends a clinician's reach is worth taking seriously. The caution is equally concrete. Almost all of this literature is built on structured electronic records from North American and European health systems, and there is no reason to expect a model fitted to those records to hold here, where much of the record is on paper or free text, where clinical documentation is often multilingual and code-switched, where case mix and pathways to care differ, and where the base rates in any given clinic differ from those in the development sample. Nothing about Indian performance is asserted anywhere in this section, because no Indian external validation of any named psychiatric model was retrieved for this chapter. Two positions hold regardless of the evidence. Accountability for a clinical decision rests with the registered practitioner and is not transferred to a vendor by the use of a tool. And identifiable patient information entered into a general-purpose service leaves the clinic, so the confidentiality obligation has to be settled before the tool is used and not after; which statutory instrument governs that transfer is a legal question this chapter has not established and therefore does not answer.

The topic reduces to four decisions.

- Establish the target before the number: what was recorded, by whom, and whether it is the construct or a proxy for it.
- Ask what the number was measured against, both in comparator and in split, and treat internal validation as an estimate of optimism rather than of transfer.
- Take the base rate into the arithmetic before accepting a screening claim, because predictive values belong to the population and not to the model.

- Keep performance, permission and patient benefit as three separate claims, each needing its own evidence.

MNEMONIC

Target, Population, Split, Comparator

The four questions that dismantle any reported figure, and the order to ask them in: what was actually being predicted, in whom, how were the evaluated data kept away from the fitting, and against what was the result judged. A claim that survives all four may then be asked the fifth and hardest: whether acting on it helps anyone.

A model's number is a measurement in one sample of one recorded label; it becomes a clinical claim only when the target, the population, the split, the comparator and the validation status are all named with it.

12 · Brain-based, biomarker-driven and computational psychiatry approaches

Psychiatry has published thousands of candidate biomarkers and holds none of them as a diagnostic test for any of its own disorders. Explaining that sentence is the item, not listing findings: grading any biological claim, from a scan to a polygenic score to a fitted model, on three questions asked in order. Can the thing be measured, does the measurement track the clinical state, and does using it leave the patient better off. Collapsing them into one is the characteristic error of the area, made as often dismissively as enthusiastically. Two neighbours are bounded out: appraising an algorithm, its dataset and its external validation belongs with the artificial intelligence and machine learning material in this chapter, and measures that have entered diagnostic use in neurodegenerative illness belong with the Dementias and neurocognitive disorders notes.

12.1 What a biomarker claim is actually claiming

A **biomarker** is a characteristic measured objectively as an indicator of a biological process or of a response to an intervention. Measured objectively excludes the rating scale, which reports experience, so a depression severity score is not a biomarker of depression however well it performs. Indicator excludes mechanism: a marker need not cause anything, and most psychiatric candidates are downstream or non-specific.

The claim must state which question the measure answers. A diagnostic marker says whether the person has the condition now, needing a reference standard and a stated comparison group. A prognostic marker says what course the person will take whatever the treatment. A **predictive** marker says which of two treatments will suit this person better, an interaction between marker and treatment. A monitoring marker says whether the state has changed, needing sensitivity to change rather than group separation. A susceptibility marker says how likely onset is in someone currently well, a statement about a population, not a person's future.

That pair is the distinction most reliably examined and most reliably muddled. A marker can be strongly prognostic and useless for choosing: knowing a patient will do badly whatever you do changes nothing. Showing a marker is predictive needs patients on different treatments compared across levels of the marker, a far more demanding design. A complete claim carries five things: the measure, the question, the population, the comparator, and the decision it feeds. Supplying the missing ones from expectation invents them.

12.2 The spine: analytical validity, clinical validity, clinical utility

Analytical validity asks whether the measurement is a good measurement, settled in the laboratory or the scanner by the properties in the table below. No part of it mentions disease, and poor analytical validity cannot be rescued downstream: the noise is inherited by every analysis built on the measure.

Clinical validity asks whether the measurement relates to the clinical state in a defined population. **Discrimination** and **calibration** are the components routinely dropped. Discrimination is the ability to rank a person with the condition above one without it, summarised by the area under a receiver operating characteristic curve. Calibration is whether a predicted probability matches observed frequency, so a one in ten risk holds about one time in ten. A model can discriminate respectably and be badly calibrated, and a badly calibrated risk is unusable at the bedside while reading well in the paper.

Clinical utility asks the only question a patient cares about: does using this test lead to better outcomes than not using it. It compares a strategy including the test against a strategy without it, carried to a patient outcome, counting the harms of testing: cost, delay, the cascade a borderline result sets off, false reassurance, and the effect of being told you carry a biological abnormality. Almost every claim made for a psychiatric biomarker is a statement about clinical validity, sometimes only analytical validity, written in the grammar of clinical utility.

QUESTION	WHAT IT ASKS	WHAT ESTABLISHES IT	HOW IT FAILS, AND THE PHRASE THAT HIDES IT
Analytical validity	Does the assay or scan measure the quantity accurately and reproducibly	Laboratory characterisation: accuracy, precision, detection limit, batch, site and scanner effects, repeat measurement	Values that shift between sites or between visits. Hidden by "a validated assay", which names nobody's validation and no population
Clinical validity	Does the measurement track the clinical state in a stated population	Sensitivity, specificity, predictive values, discrimination and calibration, in a defined sample, against a stated reference standard, replicated	Separation that shrinks once the comparison group resembles the real differential. Hidden by "significantly associated", which says a difference exists, not that a test works
Clinical utility	Does using the test improve what happens to the patient	A strategy with the test against a strategy without it, carried to a patient outcome, counting the harms of testing	No such study exists for the marker in question. Hidden by "may help guide treatment", asserting the third question from evidence for the second

RULE **Analytical validity, clinical validity and clinical utility are asked in that order, and none implies the next.** A perfectly reproducible assay may bear no relation to any clinical state, and a measure that separates two groups beyond doubt may change no decision and improve no outcome. Grade every claim by naming which of the three it has established, and say plainly that the others are unaddressed rather than assuming they follow.

12.3 Why a group difference is not a test

A group difference says one distribution's average sits away from another's; a test has to place an individual on one side of a line. A difference can be beyond statistical doubt while the distributions overlap almost completely, because a large enough sample makes a small separation reliably detectable. Reliably detectable and usefully separated are different properties, and the first is what most published biomarker findings have.

The second defeater is the base rate, and it defeats measures that do separate. Take a test right in 90 percent of people who have a condition and in **90 percent** of those who do not, applied where **1 person in 100** has it. Of 10,000 tested, 90 of the 100 with the condition test positive, and 990 of the 9,900 without it also test positive. A positive result is therefore far more often wrong than right: **fewer than 1 in 10** positives is a true positive. The prevalence did the damage, not the test, and no gain in sensitivity undoes it, because the false positives come from the enormous group without the condition.

90%	90%	1%	8%
SENSITIVITY, STIPULATED	SPECIFICITY, STIPULATED	HOW COMMON THE CONDITION IS	CHANCE A POSITIVE IS TRUE

Those four figures are a worked illustration, not a measurement of any real test. The shape is what to carry: at low prevalence positive predictive value collapses while negative predictive value becomes excellent, so a marker can be useful for ruling out and useless for ruling in.

The third defeater is **spectrum bias**. Performance is first measured in clear, established, often severe illness against healthy volunteers; the clinic asks the hard question, an early, partial, comorbid presentation against the other diagnoses on the differential. Move the comparison group to the real differential and specificity falls, sometimes to nothing, because the marker was never specific to the disorder, only to being unwell. An accuracy figure must travel with the comparison group it was earned against; without it the figure is uninterpretable rather than merely imprecise.

TRAP **Repeating a research accuracy figure as though it described the clinic.** Published accuracy is typically obtained in a balanced sample against healthy participants, in people whose diagnosis is already certain. The clinic has none of those properties: low prevalence, a comparison group that is the differential, and patients presenting early and ambiguously, which is why the test was wanted. The defensible sentence names the sample and the comparator beside the number, and states that performance in an unselected population is unknown unless measured there.

12.4 The reference standard problem, and why a research framework is not a diagnosis

The deepest difficulty is not the measure, it is the yardstick. A diagnostic biomarker is validated against a **reference standard**, and in psychiatry that standard is a clinical diagnosis made by interview against the current manuals. Two consequences follow.

The first is a ceiling: a measure judged against clinical diagnosis can at best reproduce it, its apparent accuracy bounded by the reliability of the standard rather than the biology, so where two experienced clinicians would disagree the marker is scored against a coin. The second is a filter running the wrong way: a measurement identifying a real biological grouping that cut across the manual's categories would be scored inaccurate and abandoned, because the criterion for success is agreement with the classification it contradicts. That circularity is why agreement with diagnosis is a weak result and predicting what it does not, such as course or treatment response, a strong one.

Dimensional research frameworks, which organise study around functional constructs measured across levels from genes and circuits to behaviour and self-report, and the hierarchical taxonomies derived from the covariance of symptoms, exist because of that problem: they free research from the assumption that the manual's categories are natural kinds. They are research frameworks, carrying no diagnostic threshold, no validated decision rule and no treatment recommendation attached to any dimension, and none allocates treatment better than a diagnosis made by interview.

GOOD TO REMEMBER

Phenomenology is not a placeholder waiting for biology to replace it. It remains the measurement instrument in psychiatry and the standard every candidate marker is validated against, so it cannot be displaced without a study showing better patient outcomes rather than better agreement. The viva position is neither that biology will soon supply the diagnosis nor that it offers nothing: these approaches produce constructs, stratification hypotheses and research designs, and no claim to have replaced the interview has been made good for any primary psychiatric disorder.

12.5 What is replicated, what it licenses, and what it does not

Group-level findings here are real; the discipline is to say what each licenses. Every row below has been reproduced, none is a test, and the fourth column carries the marks.

MEASURE FAMILY	WHAT REPLICATES AT GROUP LEVEL	WHAT THAT LICENSES	WHAT IT DOES NOT LICENSE
Structural imaging	Enlarged lateral ventricles in chronic schizophrenia, among the field's oldest and most reproduced findings, and small average reductions in cortical and subcortical volumes across several disorders	A statement about averages, and a research target	Individual inference. Distributions overlap heavily with comparison groups and the differences are not specific to one diagnosis
Functional imaging and connectivity	Group differences in task activation and in resting connectivity across mood, psychotic and anxiety disorders	Hypotheses about circuits, and stratification questions worth testing prospectively	A scan for a patient. Effects are small against within-person and between-scanner variability
Electrophysiology	Reduced auditory event-related potential amplitudes and reduced mismatch negativity in schizophrenia, among the most replicated findings in psychiatry	Research into information processing, and endophenotype work in relatives	Diagnosis or exclusion in an individual. The reductions are neither specific nor separating
Peripheral and inflammatory measures	Modest average elevations of circulating inflammatory markers in a subgroup of people with depression	A stratification hypothesis, to be tested rather than applied	Diagnosis, and any causal statement. Confounded by adiposity, smoking, infection and physical comorbidity, and specific to no psychiatric disorder
Neuroendocrine	Group differences in cortisol regulation, the historical example being the dexamethasone suppression test proposed for melancholic depression	A permanent teaching case: it separated groups, and specificity collapsed once the comparison group was other psychiatric patients	Diagnostic use, which was proposed, tried and abandoned
Common-variant genomics	Polygenic scores associated with case status for schizophrenia and for bipolar disorder at the group level, robustly and repeatedly	Research into shared architecture across diagnoses	Individual prediction, diagnosis or counselling. Variance explained is small, distributions overlap almost completely, transfer across ancestries is poor

The one place where fluid and imaging measures have crossed into diagnostic use is neurodegenerative illness, taught with the Dementias and neurocognitive disorders notes; nothing about their performance transfers to a primary psychiatric disorder, and citing them as

proof that psychiatry now has biomarkers is a category error. And this section makes no regulatory statement of any kind: the chapter's status ledger holds one row per drug per jurisdiction and no row for any assay, imaging measure, device or diagnostic test, so no approval, clearance, qualification or reimbursement claim appears here for any country. Demand of anyone who offers one the four fields the chapter uses throughout: the jurisdiction, the exact intended use, the authoritative source, and the date the status was true. An intended use is bound to the jurisdiction that granted it, and a test available in one country establishes nothing about another.

12.6 Computational psychiatry: what a model parameter is, and what it is not

Computational psychiatry states a hypothesis about a mechanism as a formal model, fits it to an individual's behaviour or brain data, and treats the fitted quantities as candidate measures: how a person processes information, in numbers comparable across people.

Two strands run under the name. The theory-driven strand builds a **generative model**, equations describing how choices or percepts are produced, and estimates its parameters for each person: reinforcement learning models yield learning rates and sensitivity to reward and punishment; evidence accumulation models yield an accumulation rate and a decision threshold, separating a slow response caused by cautious responding from one caused by inefficient processing; Bayesian and predictive coding models yield the weight given to prior expectation against incoming evidence, describing hallucination and delusion as an altered balance rather than an unexplained content. The data-driven strand predicts without committing to a mechanism, and its discipline is taught with the artificial intelligence and machine learning material.

A fitted parameter is an estimate produced by fitting a model you chose to data you collected, a property of the model and the data together rather than a substance inside the patient, so changing the model changes it. Four requirements follow beyond the ordinary ones. **Parameter recovery** comes first: simulate data from known values, refit, and check that the known values are recovered, because a parameter unrecoverable from data the model itself generated cannot be interpreted in a patient. Identifiability next: two mechanisms producing indistinguishable behaviour cannot be told apart by fitting. Third, model comparison is relative, selecting the best of the models written down, and the best of a bad set is still that. Fourth, prediction must be out of sample, with any selection of features or participants kept inside the training data.

Apply the spine and a computational measure reads cleanly. Analytical validity becomes parameter recovery plus the stability of the estimate when the same person repeats the task. Clinical validity becomes whether the parameter separates clinical groups or predicts an outcome. Clinical utility becomes whether allocating treatment by the parameter beats allocating it clinically. Computational psychiatry is substantial at the first two and effectively silent at the third, and why the third is a different study and not a bigger one is the examinable point.

12.7 Reliability, replication, and the claim that a biological measure is unbiased

Two mechanical facts explain most of the disappointment here. The first is a reliability trap in task design: a task is optimised to produce a robust average difference between conditions, done

by making everyone respond as similarly as possible, deliberately minimising the between-person variation an individual measure needs. So a group effect that replicates without fail can sit beside a between-session correlation close to nothing in the same individuals, both consequences of one design choice rather than a contradiction. Because reliability caps the correlation a measure can show with anything else, an unreliable measure cannot correlate strongly with symptoms or outcome however real the relationship. A measure built for group inference must be rebuilt before it serves as an individual marker.

The second is that discovery is not validation. Clustering a sample into subtypes always returns clusters, whose number and shape depend on the measures entered, the preprocessing and the algorithm. Several prominent psychiatric subtyping solutions have not been recovered in independent samples. A subtype derived in one dataset is a hypothesis until an independent dataset recovers it and it predicts something outside the data that generated it.

PITFALL

Treating poor test to retest reliability as a minor technical caveat. It is not a footnote to the finding but a limit on it. If a measure taken twice in the same stable person gives substantially different answers, it cannot support an individual clinical decision whatever its group statistics look like. Ask what the between-session agreement was before asking what the between-group difference was, and note how often that question is unanswered.

The claim that measuring biology removes human bias is false, and the four routes by which bias enters are a genuine discriminator. The label carries it first: a marker trained against clinician-assigned diagnoses learns whatever systematic patterns those clinicians produced, reproducing the bias with a more authoritative surface. The sample carries it second: performance is a property of the population studied, shown most clearly by polygenic scores, whose performance falls substantially when transferred from the ancestry group they were derived in. The instrument third: scanner, batch and site effects are real, and where sites differ in whom they recruit, a site effect and a demographic effect cannot be separated afterwards. Access fourth: who is tested determines who can benefit, and private testing is distributed by income rather than need.

IN INDIA

The transfer problem is not abstract here. Reference intervals, normative neuropsychological data and polygenic scores are largely derived from populations in which South Asians are thinly represented, so a value abnormal against an imported reference may be unremarkable for the person in front of you, and a risk score performs worst precisely where it is newest to the market. Two consequences. A group frequency is a prior about a population and never a result about a patient, so no measure may be read off ancestry. And before treating a value as abnormal, ask where the normative sample behind that threshold came from; where local normative data exist, use them.

12.8 Where a biological test does change management, and what a complete answer contains

An answer finishing on scepticism has missed the clinical half of the item. The tests that earn their place are not tests for psychiatric disorders but for conditions that both imitate a psychiatric presentation and carry a different treatment, and their utility is easy to demonstrate because a positive result changes management at once.

The categories follow from the principle rather than a memorised panel: endocrine disturbance, particularly thyroid disease; nutritional deficiency; infective and inflammatory causes; antibody-mediated encephalitis, characteristically a rapidly evolving first episode of psychosis with atypical features, seizures, movement abnormalities or autonomic instability; inherited metabolic disease such as Wilson disease in a young patient with psychiatric and movement symptoms together; structural and epileptic causes; and substance and medication effects, commoner than all the rest. Justify each test by the decision it could change, and notice that this is the one part of the topic where the third question of the spine is genuinely satisfied.

IN INDIA

Two situations recur in outpatient practice. A family arrives having paid privately for a brain mapping report, a neurotransmitter panel or a genetic wellness profile and asks what it means. The question is not whether the number sits outside its range but which decision would change if it were different, and where the honest answer is none, saying so is the correct clinical act, with an explanation of why the test could never answer the question it was sold as answering. An unvalidated test is not neutral: it has cost the family money and it generates false reassurance. The second situation is the opposite failure, the treatable mimic missed because nothing was investigated.

TRAP **Writing that a scan or a blood test can now diagnose a psychiatric disorder, or that it soon will.** Neither half is defensible. There is no validated transdiagnostic blood test or brain scan for a primary psychiatric disorder; a group-level association does not determine an individual diagnosis; and a research classification organises study rather than replacing the interview. The forecast is not safer than the claim, because a prediction about validation is a claim about evidence that does not exist. Say what has been established, name which of the three questions it answers, and state that the others are open.

Four moves answer this item in whatever form it is asked.

- Grade the claim on the spine: analytical validity, clinical validity, clinical utility, named in order, with the one actually established stated and the others called open.
- Give the two arithmetic reasons a group difference is not a test, overlapping distributions and the base rate, with spectrum bias as the reason published accuracy falls in clinic.
- Name the reference standard problem: a marker validated against clinical diagnosis is capped by it and discarded if it disagrees, so predicting course or treatment response is the stronger design.
- Treat a computational parameter as a property of the model and the data, needing recovery, identifiability, out-of-sample testing and test to retest reliability before it measures anything

in a person.

MNEMONIC

Measure, Separate, Decide, Benefit

The first three are the spine and the fourth is the study almost nobody has done: does the result change a decision, and does that change leave the patient better off. Every biomarker claim sits at the step where its evidence stops.

A biomarker is graded on three questions asked in order, a group difference is not a test, and the yardstick every candidate is judged against is still the clinical interview.

13 · Digital therapeutics and prescription apps in psychiatry

A prescription digital therapeutic is a treatment whose active ingredient is software, and almost every difficulty follows from that sentence. The clinical question is how an intervention delivered on a patient's own device is evaluated, when its effect depends on how much of it the patient does, when no convincing dummy version can be built, and when the user cannot be kept unaware of what they are using. The regulatory question is what kind of object it is, since it is not a medicine and does not reach a market by a medicine's route. Marks are lost in three places: appraising the product as though it were a drug, reading an effect measured in the patients who kept using it as the effect of the product, and supplying a regulatory status from memory because none was written down. Two neighbours share this chapter's readiness ladder and neither is rebuilt here: appraising a fitted model, its dataset and its external validation belongs with the artificial intelligence and machine learning material, and grading a measure on analytical validity, clinical validity and clinical utility belongs with the biomarker material.

13.1 What the object is, and the four things it is not

A **prescription digital therapeutic** is a software product that makes a therapeutic claim about a named condition, delivers the intervention to the patient rather than to the clinician, and is gated behind a clinician's authorisation. The claim decides the category, not the content: two products can ship identical exercises and be different objects entirely, because one asserts treatment of a named condition and the other does not.

Four neighbours are routinely collapsed into it. A wellness or self-help application may carry the same material, makes no therapeutic claim, and is therefore assessed by nobody. A telemedicine platform delivers a clinician at a distance, so the treatment remains the clinician's and the software is a channel. A decision-support tool acts on the clinician's judgment and the patient never touches it, which is the object the artificial intelligence and machine learning material appraises. Software supplied alongside a medicine carries no therapeutic claim of its own, because the claim belongs to the drug.

The sentence that does for software what a molecule does for a drug is its **intended use**: the population, the condition and the clinical purpose, written down. Everything is assessed against

it, and a product whose intended use is amended is a different product.

13.2 Why it is a device and not a drug, and the six things that follow

A medicine is a substance and its manufacture reproduces that substance. Software has no substance: what is assessed is a design together with an intended use, and what is produced is a build. Hence the category **software as a medical device**, software intended for a medical purpose that achieves it by processing information rather than by chemical or physical action on the body.

First, the unit of authorisation is one product, one intended use and one version, so there is no class of digital therapeutics of which a product is an instance, and a class-level statement of status is not the kind of statement that could be true. Second, risk classification drives the evidence burden: device frameworks grade by the harm a failure could cause, and lower grades may rest on similarity to something already lawfully marketed rather than on a fresh trial, so two products may be marketed on entirely different evidence. A permission tells you a pathway was satisfied, not what was shown. Third, there is no generic: a second manufacturer cannot show its software equivalent to another company's, so it builds a different product with its own authorisation and evidence.

Fourth, **change control** is continuous rather than exceptional: a tablet's formulation is frozen at approval, software ships updates, and since the object assessed was a version, the question after launch is whether the thing in the patient's hand is still the thing assessed. Fifth, withdrawal has no drug analogue, because a dispensed medicine stays with the patient, whereas a therapeutic depending on a running service can stop working and the patient loses the treatment with no clinical decision taken. Sixth, being authorised creates no way to deliver it: a prescription for a molecule meets a pharmacy, while a prescription for software meets whatever arrangement a manufacturer, a service and a payer have built, and where none exists it cannot be filled.

RULE **The thing authorised is one product, for one intended use, at one version; move any of the three and the permission does not travel with you.** Three errors share that shape: a claim about a category, a claim about a wider population than the intended use names, and a claim about the current release when the evidence came from an earlier one, the last invisible unless somebody asks which version was studied.

13.3 What this chapter can see, and what it structurally cannot

The most useful thing this section can give a candidate is a worked example of an absence that means nothing. This chapter carries a status ledger of one row per agent per jurisdiction, each row carrying the same four mandatory fields, the jurisdiction, the indication or a declared gap in its place, the authoritative source or that same gap, and the date the status was true. It holds **48** rows over **16** agents in three jurisdictions, every one a medicine. The number of rows describing any software product is **0**.

That is not an oversight, and reading more of the same documents would not close it. The instruments recorded as used number **7**. Five can see nothing but medicines: two interfaces onto a national drug application and labelling archive, a register of centrally authorised medicines

with its assessment pages, and a corpus of published new-drug approval lists. The sixth is a whole-domain search of the drug regulator's own website, pointed in this build only at drug names. The seventh searches the national trials registry, which records research and never marketing status. A drug application archive cannot return a device record, because that is not the kind of record it holds. The five archives are silent about software by construction, the other two by what they were pointed at.

Two facts in this chapter's working files make that sharper. The source registry that governs which document is authoritative names, for artificial intelligence and digital therapeutics specifically, a device database, and that database appears neither among the instruments recorded as used nor among those reachable but not yet run. And the grounding lanes name the drug items and do not include this item at all, not even a status lane. The absence here is not a search that came back empty; it is a question none of the records these instruments hold could answer.

48	16	0	7
ROWS IN THE CHAPTER STATUS LEDGER	AGENTS COVERED, ALL MEDICINES	ROWS FOR ANY SOFTWARE PRODUCT	INSTRUMENTS USED, NO DEVICE REGISTER

The falsification pass built to test those rows compares a stored indication text against a cached primary drug document, and a software product would not fail those instruments: having no drug label to read, it is inapplicable to them. An instrument that cannot fail is not a test.

So no product is named in this section, and no approval, clearance, authorisation, licence, marking, date or jurisdiction is asserted for anything. Two sentences are forbidden and the second is the one that feels safe: that a named product is approved somewhere, and that a named product is not approved somewhere. The first would be written from memory; the second would rest an assertion on the silence of instruments that hold no such record.

TRAP **Filling the regulatory blank yourself, in either direction.** A reader who meets a section naming no product and no status will supply both, because a refusal reads as an omission. Half-remembered regulator statements are the commonest fabrication here, and the mirror error is worse, concluding from silence that nothing has been authorised anywhere. The examinable move is to name the record that would settle it, state that such a record is per product, per intended use and per version, and say none was retrieved.

13.4 How the evidence differs from a drug's evidence

The trial machinery looks familiar and diverges at six points.

QUESTION	A MEDICINE	A PRESCRIPTION DIGITAL THERAPEUTIC
What is administered	A fixed quantity of a substance	An opportunity to do a course of work, which the patient may not take
How exposure is known	Inferred from prescriptions and self-report	Logged continuously, module by module
What varies between patients	Absorption, metabolism, adherence	How much was received at all, from none to all
What blinding is possible	Allocation and the treatment itself	Allocation only; the patient knows what they are using
Who measures the outcome	Often a rater unaware of allocation	Usually the unblinded patient, on the same device
After authorisation	A frozen formulation	Successive versions, of which the studied one may no longer be in use

The exposure row cuts both ways. Precise logging is a real advantage, since treatment received is measured rather than assumed. The same precision creates the field's most attractive error: because exposure is known and varies enormously, an analysis restricted to those who completed the programme is available, tempting and severely biased, since completers are the engaged and engagement predicts outcome for reasons unconnected to the software. The randomised comparison lives in the analysis that keeps everybody allocated to it.

13.5 Engagement and attrition, which do something a pill's adherence does not

A tablet delivers its benefit on being swallowed. A therapeutic delivers its benefit only when the patient does the work inside it, so the active ingredient is the patient's own activity. **Engagement** is therefore not adherence renamed: adherence is a gate the benefit passes through, whereas engagement is part of the benefit, which is why an effect size is uninterpretable without a statement of how much of the intervention was taken up.

Attrition then does four separable things. It empties the denominator, steeply, since disengagement is heaviest early. It is not random, so where the arms lose different sorts of patient the randomisation that made the comparison fair has quietly been undone. It is itself an outcome rather than a nuisance, because a treatment patients stop using is not working. And it is entangled with the state the product exists to change, since low motivation, poor concentration and hopelessness are features of the conditions treated and also the commonest reasons a person stops opening an application. An effect conditional on sustained use is conditional on the patient not having the problem in its severe form.

Two further effects have no drug counterpart. Novelty decays, so an effect measured across a short window need not describe the same product used for a year. And the human contact around a trial is a co-intervention: visits, reminders and a coach who telephones when use lapses deliver

support, so a therapeutic evaluated with that scaffolding is a different intervention from the software used alone, and the unsupported version is usually what a service plans to deploy.

GOOD TO REMEMBER

Prescribing one of these is prescribing a course of work, so time the review to where the work stops, not to the end of the programme. Silent abandonment is the ordinary outcome, and a patient who has quietly stopped is often assumed by everyone to be in treatment. Ask what the patient did, in what week they last did it, and what got in the way, and treat disengagement as clinical information, not non-compliance.

13.6 Why the control condition is unusually hard to build

Every comparator here fails to control something, and naming what each leaves uncontrolled is the appraisal skill worth carrying.

COMPARATOR	WHAT IT CONTROLS	WHAT IT LEAVES UNCONTROLLED
Waiting list	Time and spontaneous change	Attention, expectancy, structured activity and being monitored, so the difference is an upper bound, not an estimate
Usual care	Whatever the service delivers	What that is: unstandardised, differing between sites, and in an under-resourced service possibly very little
Attention-matched or sham application	Screen time, contact, expectancy, daily use	Whether it is truly inert, since anything engaging enough to be used is doing something, and whether patients believed it
The same treatment delivered by a person	The content, isolating the mode of delivery	Nothing important, which is why it is the informative arm, the most expensive and much the rarest

The **digital placebo** names the problem: the benefit that comes from using something at all, from the daily ritual, the attention and the expectation of improvement. A **sham comparator** must satisfy three requirements in tension: plausible enough to be believed, engaging enough not to be abandoned faster than the active arm, and inert. Anything meeting the first two risks failing the third, and differential dropout between a dull sham and an interesting therapeutic reintroduces the bias the sham was built to remove.

Blinding compounds it. Allocation can be concealed at randomisation, but a participant cannot be kept unaware of what they are using for weeks, and the outcome is usually their own self-report on the same device, so the person who knows which arm they are in is also the measuring instrument. That is a structural weakness rather than a defect of any one study, and a larger sample does not repair it. Two consequences: an outcome assessed by someone unaware of allocation, or one not self-reported at all, carries much more weight here than a self-rated scale;

and framing is part of the intervention, because content presented as prescribed and authorised generates an expectancy the same content offered as a self-help download does not.

TRAP **Reading a waiting-list-controlled result in completers as the effect of the product.** It stacks two errors working the same way. The waiting-list comparison leaves attention, expectancy and structured activity uncontrolled, so it measures the difference between doing something and doing nothing rather than the difference this product makes; restricting the analysis to completers then removes everyone in whom the product failed to hold attention. Ask what the comparator was, who was in the denominator and who assessed the outcome before accepting any figure.

13.7 What an Indian service would have to establish, and what a complete answer contains

The lens here is the product: what the evidence shows, what kind of regulated object it is, and what a prescription for it would mean. Access, equity, service integration, adoption and programme monitoring for digital mental health in this country are the subject of the Community psychiatry and public health notes, because whether a product works and whether a system can deliver it are answered from different evidence.

IN INDIA

Six things must be established before a service deploys one, and none may be inferred from an application store listing, a company's own description, or a permission granted in another country. First, what kind of regulated object the product is here and on whose authority, with the document named and dated, since no instrument in this chapter's set can return that record and nothing in this section supplies it. Second, whether the evidence was generated in a population and language resembling the patient, and whether the content was rebuilt for that language or merely translated, since a therapeutic is made of words in a way a tablet is not. Third, which version the patient will receive, and whether it is the version studied. Fourth, who holds the data it generates and where that goes, since a therapeutic collects a clinical record continuously as a condition of working. Fifth, what happens when the patient stops, and who notices. Sixth, whether the human support that surrounded the product in its trials can be supplied here, because without it the intervention deployed is not the one evaluated.

Reduced to four decisions, the item is answerable in any form.

- Identify the object first: a therapeutic claim about a named condition, delivered to the patient, gated by a clinician, assessed as one version of one product.
- Appraise it as a device trial: what the comparator controlled, who was in the denominator, who rated the outcome, how much was received.
- Treat engagement as part of the treatment and attrition as a result, and refuse any effect size quoted without both.
- Keep evidence, permission and deliverability separate, and where no regulatory record has been retrieved, name the missing record instead of guessing what it would say.

MNEMONIC

Claim, Version, Comparator, Completion

Claim is what makes software a regulated object and fixes what may be said about it. Version is what was assessed, and the question nobody asks. Comparator is what the trial controlled, and so what its number means. Completion is how many were still using it when the outcome was measured, both a result and the limit on every other result.

The active ingredient is the patient's own work, the object authorised is one version of one product for one intended use, and an absence reported by an instrument built for medicines is not a fact about software.

Individual newer agents

14 · Asenapine: pharmacology and clinical use

One active substance, two routes, and four regulatory records that do not agree on what the drug is for. Asenapine is the most compact item in this chapter and the easiest to answer wrongly with confidence: a candidate who knows it from the bipolar algorithm writes an indication correct in one country and false one border away. What follows is the whole-agent profile: the formulations, the indication set jurisdiction by jurisdiction, and the discipline that keeps the two fastened together. Receptor pharmacology, the sublingual dosing detail and asenapine's place in the bipolar phase algorithm are taught in the Mood disorders: pharmacotherapy notes.

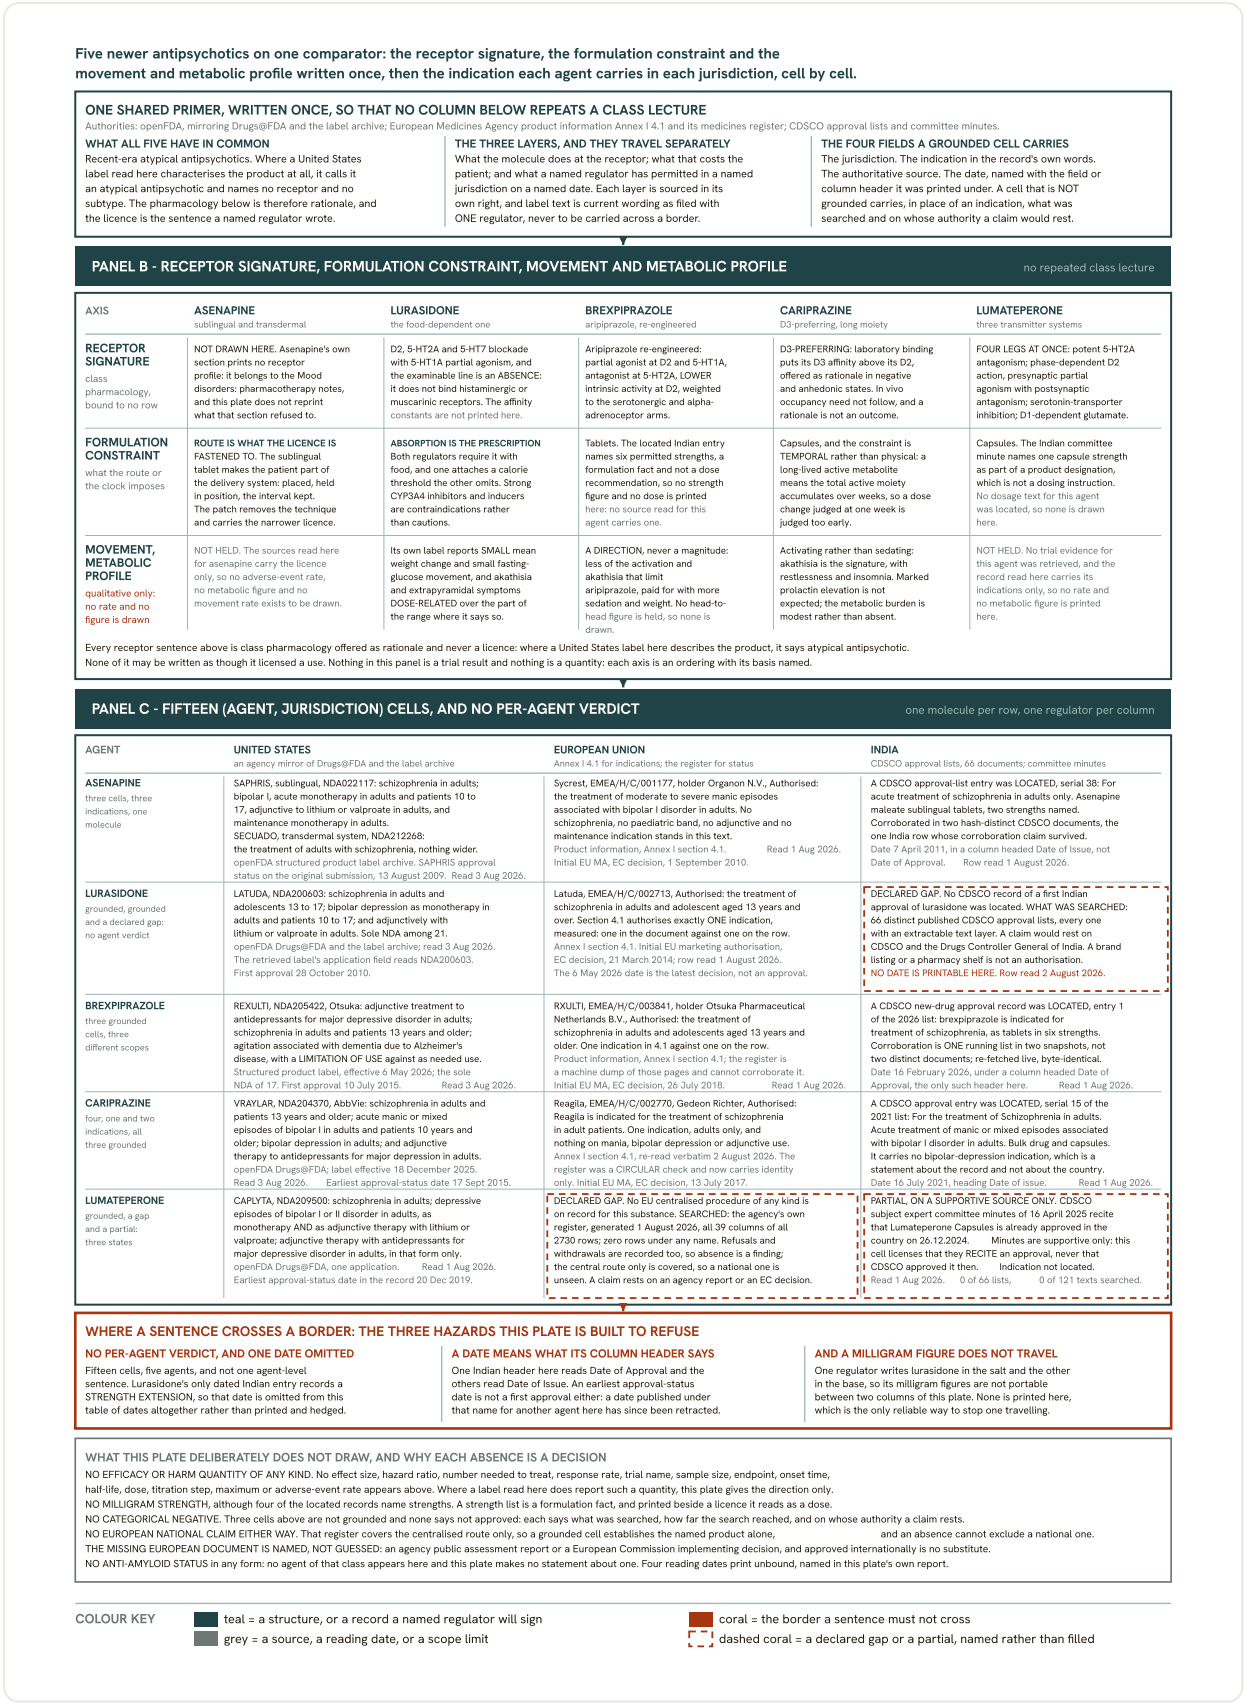


Fig 38.8 Five newer antipsychotics on one comparator: fifteen (agent, jurisdiction) cells, not five verdicts. Where a United States label read here describes the product it says only atypical antipsychotic, so the receptor row is rationale and the licence is the regulator's sentence. Movement and metabolic effects are qualitative: no rate, dose or onset time appears. The states differ inside one molecule: lurasidone's Indian record is a declared gap, lumateperone's recites an approval it does not establish, and no cell is a per-agent verdict. A cell that is not grounded names what was searched. (openFDA; European Medicines Agency Annex I 4.1 and medicines register; CDSCO lists and committee minutes.)

14.1 Two American applications, and why they are not one drug

The United States holds two applications for asenapine, and they are not two names for one product. The **innovator application** is a **sublingual tablet**, SAPHRIS: application NDA022117, sponsor ABBVIE, whose original submission carries approval status dated **13 August 2009**, in three strengths, 2.5 mg, 5 mg and 10 mg expressed as asenapine base. The second is a **transdermal system**, SECUADO: application NDA212268, sponsor HISAMITSU, whose original submission carries approval status dated **11 October 2019**, as three patches delivering 3.8 mg, 5.7 mg and 7.6 mg over 24 hours.

They do not license the same thing. The sublingual label reads: SAPHRIS is indicated for schizophrenia in adults, and for bipolar I disorder as acute monotherapy of manic or mixed episodes in adults and pediatric patients 10 to 17 years of age, as adjunctive treatment to lithium or valproate in adults, and as maintenance monotherapy treatment in adults. The transdermal label reads, whole: SECUADO is indicated for the treatment of adults with schizophrenia. One route carries schizophrenia and a three-part bipolar I indication reaching to age 10; the other adult schizophrenia alone.

Three constructs inside that indication are routinely collapsed together. **Acute monotherapy** of manic or mixed episodes is licensed in adults and in patients aged 10 to 17. **Adjunctive treatment** to lithium or valproate is licensed in adults only, and so is **maintenance monotherapy**. The paediatric band sits on one of the three and not on the other two, so an answer that calls the drug licensed in adolescents without saying for what has widened a licence by omission.

RULE **A regulatory sentence names four things or it is not one: the jurisdiction, the exact indication in the record's own words, the authoritative source, and the date the status was true.** On this agent the fourth field alone splits: the American row was read on 3 August 2026, the European and Indian rows on 1 August 2026. Copying one reading date onto the others is the same error as copying an indication across a border.

14.2 The row that blended three products, and how it was caught

The most transferable thing about this item is not the drug; it is the way its record went wrong. The first American row was built by querying the regulator's interface on the generic name, returning five applications, three of them abbreviated applications for generic products, and a set spanning SAPHRIS, SECUADO and generics. A falsification pass rejected it on two instruments. The earliest date across a mixed set is not a first approval date but the earliest among whatever the query returned; and a set spanning more than one product means the label and dates retrieved may describe different products.

What the repair found is worse than a wrong date. The row cited SECUADO, carried the date belonging to SAPHRIS, and published the sublingual bipolar text that SECUADO does not carry: citation, date and indication text each pointed at a different application. Nothing in its appearance signalled that. It now holds the innovator with its own matching date. The repair instruction is the method: query per application rather than per generic name, paginate to

exhaustion, and write one row per application with its own sponsor, brand, approval date and label, never merging applications into an agent-level row.

The row was then tested rather than trusted. Three text instruments were run against the primary label: that the row text is verbatim, that it is not the primary text cut short, and that the count of indications carried matches the document. All three passed, at 680 characters against 680 and one indication against one. It also carries a caution about itself: the interface is an agency-operated mirror of the approval database and the label archive rather than the application itself, and the wording is current label text as filed, which must never be carried across a jurisdiction.

PITFALL

Searching a regulator by drug name and reading the first date that comes back. A generic-name query returns innovator and generic applications together, and the earliest date in that heap belongs to whichever record happened to be earliest, not to the drug. For asenapine the heap held three generic applications out of five. Ask which application, then which product, then which date, in that order.

14.3 Route is what the indication is attached to

Asenapine is given sublingually because a swallowed tablet does not deliver a useful amount of drug; the pharmacokinetic reason, the dose and the timed post-dose instruction belong to the Mood disorders: pharmacotherapy notes. A sublingual tablet makes the patient part of the delivery system: the tablet has to be placed, left in position, and the interval afterwards kept, or the dose administered is not the dose prescribed. A transdermal system moves that burden onto an application site.

That is why a transdermal asenapine exists, and why the trade is not free. The route that removes the technique is the route with the narrower licence: adults with schizophrenia and nothing more, against a sublingual label that also carries bipolar I disorder. A patient with bipolar I mania who cannot manage a sublingual tablet cannot simply be moved to the patch.

One boundary must be stated. No trial-level evidence is held for this agent in this build: item 291 carries a status lane and no evidence lane, so nothing here names a pivotal trial, a sample size, an endpoint, an effect size or an adverse-event rate. The comparison above is of licences and not of efficacy.

14.4 The European Union: one product, one indication, not the American one

Europe holds one centrally authorised asenapine product: Sycrest, product number EMEA/H/C/001177, status Authorised, whose **marketing authorisation holder** is Organon N.V., with an initial European Union marketing authorisation, by European Commission decision, dated **1 September 2010** and a latest European Commission decision of 25 February 2026.

The indication, verbatim: Sycrest is indicated for the treatment of moderate to severe manic episodes associated with bipolar I disorder in adults. Read what is absent. No schizophrenia. No

paediatric age band. No adjunctive indication with lithium or valproate, and no maintenance indication. And a severity qualifier, moderate to severe, that the American text does not carry.

Where that sentence comes from is examinable in itself. Its authority is the **Summary of Product Characteristics**, Annex I section 4.1. The row originally took its indication from the agency's published medicines register, citing the assessment-report page as an independent second instrument; the register is a machine dump generated from those same pages, so the corroboration was circular and blind to the very failure it existed to catch. The register and that page now corroborate product identity and authorisation status only. When the eight grounded European rows here were re-read against the product characteristics, six differed and two were already verbatim; asenapine was one of the two. A row that turns out to be right is not a row that has been checked.

One limit travels with every European sentence here: this is an authorisation for the named product through the **centralised procedure**. It says nothing about other products containing the same active substance that may be authorised nationally through the decentralised or mutual-recognition procedures, which the register does not cover by construction. The defensible sentence is about Sycrest, not about asenapine everywhere in the Union.

TRAP **Writing the American bipolar indication into a European or an Indian answer.** The European authorisation carries moderate to severe manic episodes associated with bipolar I disorder in adults and nothing else. The Indian entry carries acute schizophrenia in adults and nothing else. Neither carries the American bipolar text. Each record is correct where it was issued and wrong where it was not, and the wrongness is invisible unless the regulator is named.

14.5 India: the one corroboration claim that survived checking

An Indian approval-list entry for asenapine was located, which makes this one of the few agents here whose Indian sentence is affirmative. The entry reads: For acute treatment of schizophrenia in adults only. Asenapine Maleate Sublingual Tablets 5mg/10mg. It sits at serial number 38 of a published CDSCO approval list, dated **7 April 2011**. Note that it names two strengths where the American sublingual product carries three.

Be exact about what that date is, because the field was corrected for saying otherwise. It is the figure printed in a column headed "Date of Issue", not "Date of Approval". The field previously claimed to quote a date of CDSCO approval as printed in the source, which asserts a label the document does not carry; it now names what is actually printed, while recording that the quantity is almost certainly the approval date, since these documents are titled lists of approved new drugs. A date is labelled with the name of the column it came out of.

This is also the only Indian row in the build whose corroboration claim survived checking. Two genuinely distinct CDSCO documents carry the entry with the identical date, verified by content hash; two other agents made the same claim and it did not survive, theirs being two snapshots of one list. Establishing it needed a rendered page rather than a text extraction. One of the two documents draws no vector rules at all, so every rule-based reader refuses it; it was read from the page rendered at 200 dots per inch, showing serial 38, its acute schizophrenia indication and the

date 07.04.11, and serial 39 dated 16.04.11 below. A neighbouring row with a different date is what tells you the reader did not drift a line.

One further control sits in those lists. Entries are tagged as an additional strength when they are strength extensions rather than new approvals, and the tag is used consistently: serials 35, 36, 37 and 41 around asenapine carry it, the asenapine entry does not. That is what makes this a located approval entry rather than a permission for another dose of something approved earlier. The same tag, read carelessly on another agent in this chapter, turned a confident Indian approval date into a retraction.

IN INDIA

Write it in this order, then stop. A CDSCO published approval-list entry for asenapine maleate sublingual tablets, 5 mg and 10 mg, was located; the indication it names is "For acute treatment of schizophrenia in adults only"; the date beside it is 7 April 2011 in a column headed "Date of Issue"; it is corroborated in two content-hash-distinct CDSCO documents; and the status was true as of 1 August 2026. Do not carry the American bipolar indication or the European manic-episode indication into that sentence. Note also the direction in which this instrument fails: no India row in this build asserts a categorical negative about a drug's Indian approval status, because the published approval-list corpus is demonstrably incomplete; each asserts instead that no CDSCO approval record was located, and names what was searched. Here the corpus did display the drug, so this row needs no such hedge.

14.6 The four records side by side

Four records, one molecule, four answers, each read on a stated date.

RECORD AND APPLICATION	PRODUCT, AS THE RECORD CARRIES IT	INDICATION, IN THE RECORD'S OWN WORDS	DATE, AND WHAT KIND OF DATE
United States, NDA022117	SAPHRIS, sublingual tablet, sponsor ABBVIE; 2.5 mg, 5 mg and 10 mg as base	Schizophrenia in adults; bipolar I disorder, acute monotherapy in adults and patients 10 to 17, adjunctive to lithium or valproate in adults, maintenance monotherapy in adults	13 August 2009, approval status on the original submission
United States, NDA212268	SECUADO, transdermal system, sponsor HISAMITSU; 3.8, 5.7 and 7.6 mg per 24 hours	The treatment of adults with schizophrenia	11 October 2019, approval status on the original submission
European Union, EMEA/H/C/001177	Sycrest, holder Organon N.V., status Authorised; no dosage form is carried by this row	The treatment of moderate to severe manic episodes associated with bipolar I disorder in adults	1 September 2010, initial European Union marketing authorisation by European Commission decision
India, CDSCO approval list, serial 38	Asenapine maleate sublingual tablets, 5 mg and 10 mg; no sponsor is carried by the entry	For acute treatment of schizophrenia in adults only	7 April 2011, printed in a column headed "Date of Issue"

680 of 680	3 of 5	2 of 2	9 of 9
CHARACTERS, US ROW AGAINST LABEL	GENERIC APPLICATIONS IN THE REJECTED SET	CDSCO DOCUMENTS CARRYING THE ENTRY	FALSIFICATION CELLS PASSED

Two things about that table are as examinable as its contents. Every one of the nine falsification cells for this agent, three text instruments across three jurisdictions, returned a pass, which is why the row is authorable at all. And the European and Indian rows each record that they were built only from their own regulator's sources, with no value from another jurisdiction consulted. Isolation is built into a record, not added afterwards by an author's promise.

TRAP **Merging the two American applications into one sentence about the drug.** It reads as the tidier answer and is a false one: it produces an agent licensed for bipolar I disorder by patch, which no record here supports. The blend has already happened once in this build's own ledger. If a stem gives you a route, answer within that route's application.

14.7 What a complete answer contains

Reduced to five sentences, the item is this.

- Two American applications, two products: a sublingual tablet carrying schizophrenia and a three-part bipolar I indication, and a transdermal system carrying adult schizophrenia alone.
- Route is not packaging; it is what the indication is fastened to, and moving a patient between routes can move them outside a licence.

- Europe authorises one product for moderate to severe manic episodes of bipolar I disorder in adults, a severity qualifier the American text does not carry.
- The located Indian entry is acute schizophrenia in adults only, dated from a column headed "Date of Issue", corroborated in two genuinely distinct documents.
- Every one of those sentences carries its jurisdiction, its indication in the record's own words, its source and the date it was read, and none of them travels.

GOOD TO REMEMBER

If the patient in front of you has bipolar I disorder and the asenapine available is the product the Indian approval entry describes, the indication being treated is not the indication that entry names. That is a decision to make and document deliberately, not one made by assuming that a licence seen in an American label follows the molecule across a border.

MNEMONIC

Application, Route, Regulator, Column

Application, because two exist here and for licensing purposes they are two drugs. Route, because it is what the indication is fastened to. Regulator, because every indication sentence belongs to one and travels nowhere. Column, because a date means what its column header says it means.

Two applications, three regulators, four indications, and not one of them may be moved to where it was not issued.

15 · Lurasidone: pharmacology, metabolic profile and indications

Lurasidone is the agent on which this chapter's central hazard lands twice, in opposite directions. Clinically it is chosen for what it does not do, and the small metabolic signal in its own label, not any novelty of mechanism, is why it is here. Evidentially it defined this chapter's grounding work: its American row published a generic manufacturer's approval date under a field name meaning first approval and was corrected only on 3 August 2026, while its Indian row is a declared gap that does not recover. Where lurasidone sits in the bipolar phase algorithm belongs to the Mood disorders: pharmacotherapy notes; what follows is the whole-agent profile.

15.1 The receptor profile, and the two absences that do the clinical work

The European Summary of Product Characteristics calls lurasidone a selective blocking agent of dopamine and monoamine effects, ATC code N05AE05, and prints its affinities: 0.994 nanomolar at the dopamine D2 receptor, 0.47 at 5-HT2A and 0.495 at 5-HT7, with partial agonism at 5-HT1A at 6.38 and blockade of the alpha-2C and alpha-2A adrenergic receptors at 10.8 and 40.7. The examinable line is what is missing: the same section records that **lurasidone does not bind to histaminergic or muscarinic receptors**.

The absent histamine binding is the usual explanation for the small weight signal and the absent muscarinic binding for the lack of anticholinergic burden, but that is an inference and should be

labelled as one: the document asserts no causal link, and it states no mechanism at all for the antidepressant effect commonly attributed to the **5-HT7 antagonism** and 5-HT1A partial agonism.

Two pharmacokinetic facts change prescribing. Lurasidone is **metabolised mainly via CYP3A4** and has an active metabolite, ID-14283, itself a CYP3A4 substrate, so parent and metabolite move together whenever the enzyme is pushed. Half-life is **20 to 40 hours** with steady state within 7 days, which makes a week the honest interval before judging a dose change. The European document also records 1.5-fold higher exposure in Asian than in Caucasian subjects in a population analysis, and its posology sets no adjustment for that: worth knowing here, and not a licensed instruction to dose lower.

15.2 The metabolic profile, and exactly what the comparison licenses

In the pooled placebo-controlled adult schizophrenia studies in the American label, mean weight change was **plus 0.43 kg on lurasidone against minus 0.02 kg on placebo**, and the proportion gaining 7 percent or more of body weight was 4.8 percent against 3.3 percent. Fasting glucose moved by no more than 2.6 milligrams per decilitre in any dose arm, and total cholesterol and triglycerides fell in every arm, placebo included.

The same section carries the figures that get misused. It reports, from two of the five studies, a weight change of plus 4.15 kg for olanzapine and plus 2.09 kg for quetiapine extended-release. The clinical studies section says why the arms were there: an **active-control arm was included to assess assay sensitivity**, to show the trial could detect a drug effect at all.

0.43	4.8	22.0	4.6
MEAN WEIGHT GAIN, KG, POOLED ADULT SCHIZOPHRENIA	PERCENT GAINING 7 PERCENT OR MORE	PERCENT WITH AKATHISIA AT 120 MG	MADRS POINTS OVER PLACEBO, MONOTHERAPY

TRAP **Quoting the olanzapine and quetiapine weight figures as a head-to-head metabolic result.** They come from arms included to assess assay sensitivity, in two of five studies, with no comparative hypothesis behind them. The defensible sentence is that lurasidone's own metabolic changes were small and that two active comparators in the same programme changed weight more, stated as what the label reports and not as a ranking it establishes. Note also its statement on duration: the effectiveness of the drug for more than 6 weeks has not been established in controlled studies, so a six-week metabolic comparison is not a maintenance claim either.

15.3 What the profile costs: akathisia, movement effects and prolactin

The American label publishes the price of that profile dose by dose. Akathisia and extrapyramidal symptoms were dose-related in the schizophrenia studies: akathisia in **5.6 percent at 20 mg, 10.7 at 40 mg, 12.3 at 80 mg and 22.0 at 120 mg** against 3.0 percent on placebo, with extrapyramidal symptoms rising over the same range from 5.6 to 22.0 percent. At 160 mg akathisia was reported by 7.4 percent, 9 of 121 patients, and the label states the dose relation only up to 120 mg. That cell rests on far fewer patients and licenses no claim that akathisia falls again at the top of the range.

Prolactin is the effect most often forgotten in an agent whose selling point is metabolic. Both documents record that **lurasidone elevates prolactin through dopamine D2 antagonism**, and the European text asks that patients be counselled on gynaecomastia, galactorrhoea, amenorrhoea and erectile dysfunction. In a young adult on long-term treatment it is the conversation that decides adherence.

PITFALL

Translating "metabolically favourable" into "well tolerated". Those are different axes and the same label carries both. The ward version is a patient started at a dose chosen for speed rather than tolerability, on an agent whose **akathisia rate** runs from 5.6 percent at the lowest arm to 22.0 percent at 120 mg, and then recorded as non-adherent when they stop.

15.4 Absorption is the prescription: food, CYP3A4 and two sets of numbers

In the European food-effect study mean Cmax rose approximately 2 to 3 times and AUC approximately 1.5 to 2 times with food compared with fasting; the American label puts the same phenomenon as an AUC increase of about 2-fold and a Cmax increase of about 3-fold, and attaches a threshold the European document omits: **take with food, at least 350 calories**. A patient taking the tablet on an empty stomach is not taking a low dose of lurasidone. They are taking an unpredictable one.

WHAT THE DOCUMENT SETS	UNITED STATES PRESCRIBING INFORMATION, LATUDA	EUROPEAN SUMMARY OF PRODUCT CHARACTERISTICS, LATUDA
Unit the doses are written in	lurasidone hydrochloride: 20, 40, 60, 80 and 120 mg tablets	lurasidone base: tablets of 18.5, 37 and 74 mg, doses stated up to 148 mg, conversion printed in the document
Food	with food, at least 350 calories	once daily together with a meal, no calorie figure
Schizophrenia in adults	start 40 mg, effective 40 to 160 mg, maximum 160 mg	start 37 mg, effective 37 to 148 mg, maximum 148 mg
Strong CYP3A4 inhibitor or inducer	contraindicated; ketoconazole and rifampin named	contraindicated; ketoconazole and rifampicin named
Moderate CYP3A4 inhibitor added	halve the dose; start 20 mg, maximum 80 mg	start 18.5 mg, maximum 74 mg

The European document states in its own words that lurasidone doses of 18.5, 37, 55.5 and 74 mg are equivalent to 20, 40, 60 and 80 of lurasidone hydrochloride, so the two columns are not a disagreement: one regulator writes the salt and the other the base. A candidate who reports a maximum of 148 mg in an American stem, or 160 mg in a European one, has carried a number across a border, which is this chapter's hazard in miniature. The magnitudes: ketoconazole raised lurasidone exposure 9-fold and the metabolite 6-fold; rifampicin lowered exposure 6-fold. Grapefruit juice is avoided in both documents.

GOOD TO REMEMBER

Carbamazepine, phenytoin and rifampicin are all on the contraindicated list, all three drugs a psychiatric patient may already be taking for a reason nobody mentions at referral. This is not a caution to manage with monitoring; both regulators write it as a contraindication. Read the whole chart, not the psychotropic column, and ask directly about anti-tubercular treatment.

15.5 Bipolar depression: the size of the effect and the edges of the claim

The American label establishes this indication on three placebo-controlled six-week studies, and the effect sizes are modest. As **monotherapy** in adults, 485 patients randomised to one of two flexible ranges, the placebo-subtracted change in Montgomery-Asberg Depression Rating Scale score at week six was **minus 4.6 points, confidence interval minus 6.9 to minus 2.3, for both dose ranges**, and the label states that the higher range of 80 to 120 mg did not provide additional efficacy on average compared with 20 to 60 mg. As **adjunctive therapy with lithium or valproate** in 340 patients who remained symptomatic, the difference was minus 3.6 points, minus 6.0 to minus 1.1.

Two boundaries are examined more often than the effect size. In the label's own words, the efficacy of the drug in the treatment of mania associated with bipolar disorder has not been established: this is a bipolar-depression agent, not a mood stabiliser across phases. And fewer than 1 percent of patients in both the drug and the placebo groups developed a manic or hypomanic episode, the figure to have ready when a stem asks about switch risk. Which patient should be offered this agent rather than another is the Mood disorders: pharmacotherapy notes' question, not this one.

The paediatric headline result is a placebo-subtracted 5.7 points on the Children's Depression Rating Scale, Revised, but the two regulators print different amounts of the same study: the European text records that the primary and key secondary endpoints were not met below 15 years of age, a placebo-adjusted minus 1.8 in 10 to 14 year olds against minus 8.6 in 15 to 17 year olds, and no such age split appears anywhere in the cached American label. The regulator that authorised the paediatric indication printed less of the age detail than the one that did not, so an answer that cites only the pooled figure has taken the more permissive reading without saying so.

RULE **For lurasidone the dose that helps and the dose that harms move in opposite directions, and the label says so in two places.** The higher bipolar-depression range gave no additional efficacy on average over the lower one, while akathisia and extrapyramidal symptoms rose with dose across the schizophrenia range. Escalating because a patient has not yet responded buys a documented rise in movement adverse effects against an effect the manufacturer's own study did not find. Wait the week steady state takes before concluding a dose is too low.

15.6 Three regulators, three different sentences, one that cannot be written

The most examinable fact about this agent is that its indication is not the same in two jurisdictions and is not established from any record located in a third. The American label indicates it for schizophrenia in adult and adolescent patients aged 13 to 17, for bipolar depression as monotherapy in adult and paediatric patients aged 10 to 17, and for bipolar

depression as adjunctive treatment with lithium or valproate in adults. Section 4.1 of the European Summary of Product Characteristics authorises exactly one indication, the treatment of schizophrenia in adults and adolescent aged 13 years and over. That count is measured: the multiplicity instrument found one indication in the European primary document and one carried on the row, and the truncation instrument compared the row's 103 characters against the document's 103.

JURISDICTION	WHAT THE RECORD ESTABLISHES	AUTHORITY READ	STATUS TRUE AS OF
United States	First approval 28 October 2010: application NDA200603, brand LATUDA, original submission 1, status AP, sponsor Sunovion. The only new drug application among 21 applications for this ingredient; the other 20 are abbreviated applications	Drugs@FDA and the label archive via openFDA; the retrieved label's own application-number field reads exactly NDA200603	3 August 2026
European Union	Authorised. Initial marketing authorisation by European Commission decision on 21 March 2014, product number EMEA/H/C/002713. The register's latest Commission decision here is dated 6 May 2026 and the record does not say what it decided	Annex I section 4.1 of the Summary of Product Characteristics for the indication; the EMA register and the EPAR page for identity and status	1 August 2026, indication re-read 2 August 2026
India	No CDSCO record of a first Indian approval of lurasidone was located. What was located is a strength extension. No approval date is printable in this row	CDSCO published new-drug approval lists, 66 distinct documents, every one with an extractable text layer	2 August 2026

The European row was itself repaired on 2 August 2026: the ledger had carried the indication as adults aged 18 years and over, taken from the EMA register while citing the EPAR page as an independent check, and the register is generated from those same pages, so the corroboration was circular. And the 2026 date is not an authorisation date: that column carries the latest Commission decision, proven on this very medicine, where it reads 2026 against a 2014 authorisation. The row is also bounded by procedure: it establishes a centralised European authorisation for the named product and says nothing about products of the same substance authorised nationally, which the register does not cover by construction.

15.7 The row that was wrong, and how a generic's date wore a first approval's name

Until 3 August 2026 the American lurasidone row was false in four ways at once, each a mistake anyone can make with the same public database. Its citation listed ten abbreviated applications and no new drug application at all. Its date was 3 January 2019, a generic manufacturer's approval date, held under a field name meaning first approval, against a true first approval of 28 October 2010, an error of more than eight years. Its indication text was the label of a generic lurasidone hydrochloride tablet and not of LATUDA. And that text had been cut at 900 characters by a fixed-width slice in the grounding script.

The application endpoint paginates at whatever limit is passed, which cut the result set to ten abbreviated applications with the innovator's application absent, and the label endpoint returns a single label for a generic-name query with no guarantee that it is the innovator's. Nothing in the returned text announced the substitution, because **a generic's indication section is a near-copy of the innovator's**. Reading the indication text could never have caught it. What did was checking, for every retrieved label, whether the application number the record itself carries matched the application cited. The repaired row now passes all three falsification instruments, the truncation check reporting 974 characters on the row against 974 in the primary document; the Indian row is exempt from all three, because a declared gap carries no indication text and is a claim about a search rather than about a product.

15.8 India: a declared gap, what it rests on, and what would close it

The Indian position is a gap that is a finding rather than an absence of work. Serial number 65 of the list of drugs approved from the SND Division from 01.01.2020 to 31.12.2020 reads, in its drug-name column, "Lurasidone Hydrochloride Tablets 20mg" followed by the qualifier " (Additional Strength)", with the indication "For treatment of patients with Schizophrenia" and a date of approval of 24-12-2020. The qualifier is the entire finding. An additional-strength permission records the extension of a permission that already exists, and therefore entails a prior CDSCO approval of lurasidone at some other strength, while no such earlier entry exists anywhere in the 66-document corpus. The only date this corpus offers records a strength extension, and it must never appear in a table of approval dates. Nor can the silence be read the other way, as evidence that the extension was itself the first permission, because the source document proves an earlier one exists outside the corpus. The row was demoted from grounded on 2 August 2026 and does not recover on further retrieval of this corpus.

TRAP **Writing a categorical negative about Indian approval.** No instrument used here supports a categorical negative, and the search behind the gap is measurably incomplete: independent sensitivity in the 2024 to 2026 window is zero of one, and CDSCO's own committee minutes of 16 April 2025 record a lumateperone approval dated 26 December 2024 that appears in none of the 66 lists, although the list declaring coverage of that whole calendar year carries an entry dated 30 December 2024. Absence from this corpus is evidence, and never proof. The correct sentence names the search: no CDSCO record of a first Indian approval was located across 66 published approval lists.

IN INDIA

What would close the gap is named on the row and nothing else will do: an affirmative CDSCO or DCGI permission, a current CDSCO-approved label, or an official approval-list entry recording a first approval at any strength. Subject expert committee minutes, clinical-trial no-objection certificates, CTRI registration, brand listings, pharmacy availability and manufacturer copy are none of them marketing authorisation. That last list catches people: a resident meets lurasidone as a strip in a pharmacy long before meeting it in any gazette, and availability feels like proof. Write the search rather than the impression, and where a unit already prescribes the drug, the document it should produce is the approved label it prescribes against.

Reduced to four sentences, the agent is answerable in whatever form the question takes.

- It is chosen for the metabolic profile and paid for in akathisia, and the two are dose-related in opposite directions.
- Food is not optional: one label sets a calorie threshold, the other does not, and both make the dose meaningless without a meal.
- CYP3A4 is the interaction axis, and strong inhibitors and inducers are contraindications rather than cautions.
- The indication set is American and broad, European and narrow, and Indian and not established from any record located.

MNEMONIC

Food, Enzyme, Dose, Salt

The four things that go wrong with lurasidone: given without a meal, given with a strong CYP3A4 inhibitor or inducer, escalated past the dose its own evidence supports, and its milligrams quoted from the wrong regulator's arithmetic.

Lurasidone's dose, its indication and its approval date each mean something different in each of the three jurisdictions, and the Indian date is the one no document read here contains.

16 · Brexpiprazole: pharmacology and indications

This is the agent on which a candidate is likeliest to write a true sentence about the wrong country. Brexpiprazole is a dopamine partial agonist whose licensed scope differs at every border read for this chapter: three indications in one jurisdiction, one in each of the other two. The marks separate on pharmacology, on jurisdiction, and on the comparison with cariprazine. Its position in the depression-augmentation ladder, beside lithium, thyroid hormone and the other atypical agents, belongs to the Mood disorders: pharmacotherapy notes.

16.1 Partial agonism is a claim about intrinsic activity

The word carrying this agent is **partial agonist**, and it states **intrinsic activity** rather than affinity. A partial agonist binds the receptor and produces a submaximal response, so what it does

depends on the tone it meets. Where dopamine transmission is high it competes with the endogenous full agonist and lowers signalling; where transmission is low it supplies a floor of activity instead of removing the little that was there. That asymmetry is the rationale of the class: damp the mesolimbic excess without emptying the nigrostriatal and tuberoinfundibular pathways to the same degree.

Two consequences follow. The burdens of full D2 blockade, extrapyramidal signs and prolactin elevation, are expected to be lighter rather than absent, because partial agonism reduces the depth of blockade and does not abolish it. And the class carries a signature adverse effect of its own rather than the antagonists': activation, inner restlessness and **akathisia**, which is residual agonist drive at D2 seen at the bedside. A candidate who writes that partial agonists are simply better tolerated has stated the first and lost the second.

16.2 Where brexpiprazole sits between aripiprazole and cariprazine

Hold brexpiprazole as a deliberate re-engineering of aripiprazole rather than a new idea: a partial agonist at D2 and at 5-HT1A and an antagonist at 5-HT2A, carrying lower intrinsic activity at D2 than aripiprazole and weighted more heavily toward its serotonergic and alpha-adrenoceptor arms. The proprietary class phrase attached to it, serotonin-dopamine activity modulator, is manufacturer language and not a regulatory category; describe the profile and leave the phrase alone.

The re-engineering exists for a clinical trade, and the trade is the answer: less of the activation and akathisia that limit aripiprazole, paid for with more sedation and more weight gain, because the serotonergic and adrenergic arms do more of the work. Hold that as a direction and not a magnitude: no head-to-head figure is printed here, because no source read for this fragment carries one, and a remembered effect size is worth less than a correctly stated direction.

Cariprazine separates from both on two features: it is D3-preferring at clinically relevant occupancy rather than merely D3-active, and it carries a long-acting active metabolite, so its offset after a dose change is far slower than the parent compound suggests. An adverse effect does not resolve on the timescale a resident expects, and a titration decision expresses itself over weeks. Brexpiprazole carries no comparable metabolite story, which is why examiners ask about the two together.

RULE **The mechanism is shared, the profile is not, and the licence is neither.** Three layers travel with a partial agonist: what it does at the receptor, what that costs the patient, and what a named regulator has permitted in a named jurisdiction on a named date. Each layer is sourced in its own right, and a sentence that slides from one to the next has stopped being checkable.

16.3 The United States licence, read whole, including the limitation

The American label is the widest of the three. It carries adjunctive treatment to antidepressants for major depressive disorder in adults; treatment of schizophrenia in adults and pediatric patients ages 13 years and older; and treatment of agitation associated with dementia due to Alzheimer's disease. Three indications, three populations, one molecule.

Then the clause a truncated reading loses. The label attaches a **Limitation of Use**: REXULTI is not indicated as an as needed, "prn", treatment for agitation associated with dementia due to Alzheimer's disease. Read what that limits. It does not withdraw the dementia-agitation indication; it forbids the way a busy ward is most tempted to use it, as a rescue dose written up in advance. The licensed use is scheduled treatment of a syndrome, not a chemical response to an incident.

The status was read from the openFDA structured product label archive, brand REXULTI, label effective 6 May 2026, whose application-number field returns NDA205422, the only new drug application among 17 applications for this active ingredient, the other 16 being abbreviated applications. First approval is recorded as NDA205422, sponsor Otsuka, original submission 1 carrying an approved status dated **10 July 2015**, and the reading was status true as of 3 August 2026. Jurisdiction, indication in the source's words, named application, dated reading: four fields, and an answer missing any of them is an assertion rather than a record.

TRAP **Quoting the three American indications as though they were the drug's indications.** They belong to one label in one country; the other two jurisdictions read here each carry exactly one. The failure is not ignorance of the other licences; it is the silent generalisation, writing "brexpiprazole is indicated for agitation in Alzheimer's disease" with no country attached. The label text is the current wording as filed with one regulator, and it is not a statement about any other jurisdiction and must never be carried across one. Attach the country to every indication sentence and the trap cannot close.

16.4 Europe: the same molecule, one indication

The European licence is narrower, and short enough to memorise. Section 4.1 of the Summary of Product Characteristics reads that RXULTI is indicated for the treatment of schizophrenia in adults and adolescents aged 13 years and older, and the falsification pass counted one indication in that section against one on the row. No adjunctive-depression indication appears in it. The sentence to write is that the European authorisation read here covers schizophrenia and names nothing else, which is a statement about a document rather than a claim about what exists elsewhere.

The identity fields matter as much as the text. The product is Rxulti, European Medicines Agency product number EMEA/H/C/003841, with Otsuka Pharmaceutical Netherlands B.V. as **marketing authorisation holder**, medicine status Authorised; the initial European marketing authorisation, by European Commission decision, is dated **26 July 2018**, and the most recent Commission decision on it is dated 9 July 2026. The reading was true as of 1 August 2026.

Two disciplines sit behind that paragraph. The first is which document is the authority: the indication text is taken from Annex I section 4.1 of the product information, because the agency's published medicines register is a machine dump generated from the same assessment-report pages and cannot corroborate them independently. Two views of one source are one source. The second is scope. An entry in that register establishes authorisation for the named product through the **centralised procedure**, and says nothing about other products containing the same

active substance that may be authorised nationally through decentralised or mutual-recognition routes, which the register does not cover by construction.

16.5 India: the row whose column header genuinely says approval

The Indian entry is the recent one, and it is the one an Indian examination cares about. The Central Drugs Standard Control Organisation entry read here states that brexpiprazole is indicated for treatment of schizophrenia, as brexpiprazole tablets **0.25, 0.5, 1, 2, 3 and 4 mg**. Those are permitted strengths and not a dose recommendation; no prescribing dose for this agent is printed here, because no source read for this fragment carries one, and the augmentation dose belongs with the ladder in the Mood disorders: pharmacotherapy notes.

The date and its label are the transferable part. The entry is dated **16 February 2026**, printed in a column headed **Date of Approval**, which makes this the only grounded Indian row in the chapter whose source header genuinely says approval rather than issue. Elsewhere the honest field name is the printed one: a date under a column headed date of issue must not be relabelled a date of approval by whoever quotes it. The provenance is entry serial number 1 in the list of new drugs approved in the year 2026 to date, within a corpus of 66 distinct published approval documents, with a separate bulk-drug permission for the substance dated 24 February 2026.

How strong is that row? Stronger than most. Its corroboration is one running list captured in two snapshots and not two distinct documents, a description corrected on 2 August 2026 from the stronger claim; being about five months old at retrieval it is the most recency-sensitive row in the Indian column, so it was re-fetched live and the source list came back byte-identical, which is the best evidence this row has. Two snapshots of one list is weaker corroboration than two independent documents, and saying so is what makes the row trustworthy rather than what weakens it.

TRAP **Turning a located approval into a categorical negative.** The Indian entry names schizophrenia, and the temptation is to write that brexpiprazole is not approved in India for anything else. Do not write that; nothing here supports it. The published approval-list corpus is demonstrably incomplete, so an absence in it is evidence and never proof, and the defensible sentence names what was located and what was searched. The mirror error runs the other way: a brand listing, a pharmacy shelf or a manufacturer's page is not a regulatory approval and must never be quoted as one.

IN INDIA

Write it in this order. A CDSCO new-drug approval record for brexpiprazole was located, dated 16 February 2026, under a column headed Date of Approval, naming schizophrenia and six tablet strengths, with a separate bulk-drug permission a week later. That entry establishes what it says and nothing wider, and it is not an availability or a price claim, neither of which any source read here establishes. For a patient or a family: an approval on a national list is not the same as a drug on the nearest shelf at a price the household can meet, and that has to be checked locally.

16.6 Adjunctive depression, and the comparison with cariprazine

This is where the agent earns its place, and where the cariprazine comparison stops being trivia. In the United States both drugs carry it: brexpiprazole's label reads adjunctive treatment to antidepressants for major depressive disorder in adults, and cariprazine's reads **adjunctive therapy to antidepressants** for the treatment of major depressive disorder in adults. Two dopamine partial agonists, one indication, so the choice between them is made on profile rather than on licence.

Cross the Atlantic and the symmetry breaks. The European product information for RXULTI names schizophrenia in adults and adolescents aged 13 years and older, and cariprazine's names the treatment of schizophrenia in adult patients. Neither European licence read here carries the adjunctive-depression indication, and brexpiprazole's is the one extending below adult age. An answer that reproduces the American adjunctive indication for a European or Indian patient has made the commonest error on this topic.

3	1	1 of 17	9 of 9
INDICATIONS, UNITED STATES LABEL	INDICATIONS, EUROPEAN SECTION 4.1	US APPLICATIONS THAT ARE AN NDA	FALSIFICATION CELLS PASSED

JURISDICTION	INDICATION, IN THE SOURCE'S WORDS	SOURCE AND DATE
United States	Adjunctive treatment to antidepressants for major depressive disorder in adults; schizophrenia in adults and pediatric patients ages 13 years and older; agitation associated with dementia due to Alzheimer's disease, plus a Limitation of Use against as-needed dosing for the dementia indication	Structured product label, application NDA205422, effective 6 May 2026; true as of 3 August 2026
European Union	The treatment of schizophrenia in adults and adolescents aged 13 years and older	Product information Annex I section 4.1, EMEA/H/C/003841; true as of 1 August 2026
India	Treatment of schizophrenia, tablets 0.25, 0.5, 1, 2, 3 and 4 mg	CDSCO published approval list, entry 1, 2026 list; true as of 1 August 2026

16.7 How a status sentence is checked, and the defect on this very row

This row was tested rather than trusted. Three text instruments were run against the primary document for each of the three jurisdictions and all nine cells passed: that the row text is verbatim from and attributed to the primary document; that it is not the primary text cut short; and that the count of indications carried matches the count in the document. The American cell measured 985 characters against 985 and five indications against five, the European 105 against 105 and one against one, and the Indian cell was carried in full. The Indian cell was located

geometrically, at entry 1 on page 1 within a named indication column, with two independent readers agreeing with the drawn table rules.

Now the defect this row carried, because its shape is the lesson. Until 3 August 2026 the American entry cited an abbreviated application, a generic's application, ahead of the innovator NDA205422, and its indication text had been truncated at 900 characters in mid-word by the script that built it. The date was never wrong. Sit with that. A right number under a wrong citation is more dangerous than a wrong number, because nothing about it looks incorrect and every downstream reader inherits a claim its stated evidence does not support. The check that catches it is not a check on the number: it asks whether the cited document could have established the fact, and an **abbreviated application** is a copy of an innovator's licence, never the origin of a first approval.

PITFALL

Completing a sentence that stops in the wrong place. The truncation above cut a label mid-word; its sibling defect cuts text so that it ends in a clean full stop the source never had, and nothing on the page then signals the loss. Both are invisible to a reader willing to finish the sentence from memory. When a quoted indication reads as though a clause is missing, say so and go back to the primary document. Every repair in this chapter's American column began with someone refusing to fill a gap and saying it aloud.

16.8 What a complete answer contains

Reduced to five moves, the item is answerable in any form.

- State partial agonism as intrinsic activity, then what it buys, lighter blockade burden, and what it costs, activation and akathisia.
- Separate the three partial agonists by discriminator: intrinsic activity at D2, the serotonergic weighting, and the long active metabolite that is cariprazine's and not brexpiprazole's.
- Attach a country, a source and a date to every indication sentence, and never carry one across a border in either direction.
- Quote the American limitation on as-needed use for dementia agitation, the clause most often lost and the one that changes ward practice.
- State an Indian record as what was located and where, name the printed column header, and refuse the categorical negative no instrument here supports.

GOOD TO REMEMBER

The commonest real-world error here is not choosing the wrong partial agonist. It is writing a rescue prescription for agitation in dementia against a licence that specifically excludes as-needed use, and then defending it with an indication that does exist. The indication and the manner of use are two separate permissions, and the second is where this label draws its line.

MNEMONIC

Tone, Trade, Territory

Tone: partial agonism acts on what the receptor is already seeing. Trade: less activation, more sedation and weight. Territory: three indications in one country, one in each of the other two, each with its own source and date. An indication sentence with no territory in it is not yet an answer.

Partial agonism is intrinsic activity, the trade is akathisia against weight, and no indication is true until a country, a source and a date are attached to it.

17 · Cariprazine: pharmacology and indications

One molecule, three national records, and four, one and two indications. The pharmacological question is what a D3-preferring partial agonist does that a full antagonist cannot, and what a long-lived active metabolite does to every decision after a dose change. The regulatory question is why it carries four licensed indications in one jurisdiction, one in the second and two in the third, all three grounded and each read on its own day. Marks are lost in three places: candidates upgrade a receptor rationale into a licensed claim, they reason about dose changes on a timescale this drug does not obey, and they carry the bipolar-depression indication across a border where no record read here places it. Phase matching, mixed states and where this agent sits in a bipolar algorithm belong to the Mood disorders: pharmacotherapy notes and are not rebuilt here.

17.1 Partial agonism, and the D3 preference that is the design claim

Begin with what the regulator wrote. The current United States label calls the product an atypical antipsychotic and names no receptor and no subtype. Every receptor sentence below is therefore class pharmacology taught as rationale, carried by no claim row in this build, and must never be written as though a regulator endorsed it.

A **partial agonist** occupies the receptor but produces a submaximal response even at full occupancy, and its intrinsic activity decides the net effect: at high dopamine tone it replaces a stronger signal with a weaker one and behaves as an antagonist, at low tone it supplies signal and behaves as an agonist. That is why marked prolactin elevation is not the expected consequence of adequate occupancy, and why the characteristic adverse effects are activating rather than sedating. It also explains a prescribing hazard: adding a partial agonist to a full antagonist is not additive blockade, because the newcomer competes for occupied receptors and substitutes its weaker signal, so a stable patient can destabilise during cross-titration.

The D3 claim distinguishes this agent from the other dopamine partial agonists here, and is most often overstated. Cariprazine is **D3-preferring**: in laboratory binding work its affinity for D3 exceeds that for D2, and D3 receptors sit in limbic territory tied to motivation and reward, hence the rationale in negative and anhedonic presentations. Two cautions belong beside it. Affinity in vitro need not become preferential occupancy in a living brain, where endogenous dopamine itself binds D3 far more potently than D2. And a rationale is not an outcome: no trial-level

evidence lane exists for this agent in this build, so no negative-symptom result, effect size or comparator is printed here.

RULE **A receptor preference is a hypothesis about mechanism; an indication is a sentence a regulator wrote.** They carry entirely different warrants. Write the pharmacology as rationale, write the indication in the regulator's own words with jurisdiction, source and date attached, and never let the first widen the second.

17.2 The long metabolite, and why this drug's clock runs in weeks

Cariprazine is metabolised to two **active metabolite** species, and the didesmethyl form is eliminated far more slowly than the parent. The relevant quantity is therefore the total active moiety, which accumulates towards steady state over weeks, not the two or three days a clinician's reflex assumes. No half-life figure is printed, because no claim row in this build carries one and a number written from memory is the defect this chapter exists to teach against. The shape of the curve is the examinable part.

Four consequences follow, each a clinical decision rather than a pharmacokinetic fact. A dose increase judged at one week is judged too early, its plasma consequence still rising. A reduction made for an adverse effect will not relieve it promptly, so a clinician seeing no change within days is tempted into a second reduction that compounds the first. Discontinuation does not produce a drug-free patient on the timescale the chart implies, which matters for a planned switch, for pregnancy planning, and for a new adverse effect a fortnight after stopping. And washout arithmetic learned on a short-acting agent gives the wrong answer.

GOOD TO REMEMBER

Write the review interval to the drug's clock rather than the clinic's. Where the active moiety is still accumulating, an unchanged picture at one week is not treatment failure and a new adverse effect at six weeks is not a new drug. Make one change and wait longer than feels natural: the commonest error is a rapid sequence of adjustments that leaves the eventual picture uninterpretable.

17.3 What the mechanism predicts, and the figures this build does not hold

The adverse-effect profile is predictable from partial agonism. Activating effects dominate: **akathisia** is the signature, restlessness and insomnia accompany it, both dose-related. Other extrapyramidal features occur but do not lead, and nausea and dyspepsia are common early. Because adequate occupancy is reached without complete blockade, marked prolactin elevation is not expected, and the metabolic burden is described as modest rather than absent.

Akathisia is where the two sections above meet: most likely to bring a patient back, most likely to be misread as agitation belonging to the illness, slowest to answer a dose reduction because of the long active moiety. A patient told the restlessness will settle in a day or two has been given a promise the pharmacokinetics cannot keep.

Now the boundary. Every sentence in this section is receptor-level and class-level reasoning. No adverse-event rate, no comparative metabolic figure, no weight-change number and no

comparative akathisia incidence is held in this build for this agent, because no evidence lane was grounded and the retrieved label section covers indications only.

PITFALL

Quoting a comparative metabolic or akathisia figure from memory. It sounds like something a monograph carries, which is why the fabricated version passes unnoticed. Two defensible sentences exist: the mechanism predicts an activating rather than a metabolic profile, and the size of that difference is a question for a named trial. Supplying a percentage turns a plausible expectation into a fabricated result.

17.4 The United States record, read field by field

This row carries the widest indication set, so it is the one material leaks from. The label indicates the product for the Treatment of schizophrenia in adults and pediatric patients 13 years of age and older; for Acute treatment of manic or mixed episodes associated with bipolar I disorder in adults and pediatric patients 10 years of age and older; for Treatment of depressive episodes associated with bipolar I disorder (bipolar depression) in adults; and as Adjunctive therapy to antidepressants for the treatment of major depressive disorder (MDD) in adults. Four indications, two reaching below adulthood at different ages, the third the bipolar-depression licence distinguishing this agent among the dopamine partial agonists.

The source is named and narrow: the openFDA Drugs@FDA interface, application NDA204370, brand VRAYLAR, sponsor AbbVie, with the current label effective 18 December 2025. The document is a **structured product label**, and the status true as of date on this row is 3 August 2026. The row cautions about its own instrument: openFDA is an agency-operated mirror of Drugs@FDA and the label archive rather than the Drugs@FDA application itself, and the indication wording is current label text as filed, which is not a statement about any other jurisdiction and must never be carried across one.

Then the date field, where this ledger has been burned. The earliest approval-status date in the record is **17 September 2015**, named for what it holds, not what a reader wants it to mean. It also shows a single application returned for this agent, thirteen approval-status submissions within it, and the most recent dated 18 December 2025, the label date quoted above. Whether the earliest of thirteen is a first marketing approval, the record does not say; elsewhere in this build a generic's approval date was published under a first-approval field name and retracted.

4	1 of 1	1125 of 1125	13
INDICATIONS ON THE CURRENT US LABEL	APPLICATIONS RETURNED FOR THIS AGENT	CHARACTERS, ROW AGAINST PRIMARY	APPROVAL-STATUS SUBMISSIONS IN THE RECORD

This row was tested, not trusted, and it had a real defect. Until 3 August 2026 its indication text was cut at 900 characters, mid-word, by a fixed slice in the grounding script that left no marker; the cited application was correct throughout, so the row was truncated only and never misattributed. The repaired row then passed three instruments against the primary document: the row text is verbatim, it is not the primary text cut short, at 1125 characters against 1125, and

the count of indications carried matches the count in the document at five against five. That last figure is its parse of a label printing the list twice, in the numbered section and the highlights summary; the teaching is that the counts agree, not that five things are licensed.

17.5 The European Union: one indication, and a corroboration that was circular

The European record is grounded and narrow. The authorised indication, in full: Reagila is indicated for the treatment of schizophrenia in adult patients. One indication, adults only, nothing on mania, bipolar depression or adjunctive use. The instrument: Annex I Summary of Product Characteristics, section 4.1 Therapeutic indications, read from the cached product information and re-read verbatim on 2 August 2026.

Product identity and status come from a separate place and corroborate identity only: the agency medicines register lists Reagila, product number EMEA/H/C/002770, status Authorised, and an independent recheck of the assessment-report page returned the same product number, status and authorisation date. The **marketing authorisation holder** is Gedeon Richter. The initial marketing authorisation, a European Commission decision, is dated **13 July 2017**; the most recent Commission decision on the file is dated 8 January 2026.

One correction here is worth more than the dates above. The indication was originally taken from the medicines register, with the assessment-report page cited as an independent second instrument. The register is generated from those same pages, so the corroboration was circular and blind to the very failure it was meant to catch, a stale summary layer. Section 4.1 of the annex is the genuinely independent instrument and is now the authority for indication text. Two renderings of one source are one source; citing a summary and its own summary as agreement checks nothing.

The scope limit travels with it. This establishes marketing authorisation for the named product through the European Union **centralised procedure**, and says nothing about other products containing the same active substance that may be authorised nationally through the decentralised or mutual-recognition procedures, which this register does not cover by construction.

17.6 India: the located CDSCO entry, and the two indications it carries

The Indian row is the one an Indian examination cares about, and here it is affirmative rather than absent. A CDSCO published approval list carries the entry, whose cell reads whole: For the treatment of Schizophrenia in adults. Acute treatment of manic or mixed episodes associated with bipolar I disorder in adults. Cariprazine hydrochloride bulk and Cariprazine capsules 1.5mg/3mg/4.5mg and 6mg. Two indications, both adult, and a product cell naming the **bulk drug** and the capsules at **1.5mg/3mg/4.5mg and 6mg**. Those strengths are a formulation fact from a name cell, not a dose; no dose, titration or schedule for this agent is held in this build.

The locator is exact, the date precisely labelled. The entry is serial number 15 in the 2021 CDSCO approval list, one document in the 66-document published corpus retrieved on 1 August 2026, and the date carried is **16 July 2021**. It is not called an approval date, and the reason is instructive: the field records the date printed in the column headed **Date of issue**, after an earlier version claimed to quote a label the document does not carry. It is almost certainly the approval

date, these being titled lists of approved new drugs, but the field now names what is printed. Corroboration is one CDSCO document, the date confirmed against the raw column layout rather than the flattened text line.

Now the discipline governing what the entry omits. It names schizophrenia and acute manic or mixed episodes, and neither bipolar depression nor adjunctive use in major depressive disorder. That is a statement about the record, not the country. The published approval-list corpus is demonstrably incomplete: every drug it is known to find sits in a document dated 2023 or earlier, its measured sensitivity in the 2024 to 2026 window is nil, and its positive-control tally may not warrant a gap on any agent approved after 2023. No India row in this build asserts that a drug is not approved in India; each asserts that no CDSCO approval record was located, and names what was searched. The same restraint applies to an indication absent from a located entry.

IN INDIA

Write it in this order. A CDSCO approval entry for cariprazine was located, dated 16 July 2021 under the printed heading Date of issue, covering the bulk drug and four capsule strengths, with two adult indications: schizophrenia, and acute manic or mixed episodes of bipolar I disorder. It carries no bipolar-depression indication, a statement about the retrieved record and not about Indian practice. Do not write that any use is unapproved in India: the corpus that would have to be complete for that sentence to be safe is demonstrably not.

TRAP **Carrying the bipolar-depression indication across a border.** It belongs to one record only: the American label carries it, the European annex carries schizophrenia in adults alone, the Indian entry schizophrenia and acute mania. Writing that cariprazine is indicated for bipolar depression, without naming the jurisdiction, is true in one place, unsupported in the second and beyond the located record in the third. The paediatric ages travel no better: 13 and 10 years are American label figures appearing in neither other record.

17.7 How that Indian cell was certified, and what a certified cell is

The entry sits in a scanned document, so reading a cell out of it is an extraction problem before a clinical one. Of three geometric readers in this build, two could not be used here: one refuses the merged block spanning serial numbers 14 and 15, the other is refused across the whole document because its declared page-level invariant is falsified there. The cell was certified instead by agreement between the reader working from the drawn table rules and one rendering the page as pixels and reading it by eye. Two readers of genuinely different kinds agreeing is what certification means; two text readers agreeing is one reading counted twice.

Certification earned its cost at once. The moment the cell became certifiable, the completeness instrument failed the row: the published text had dropped the bulk-drug leg of the name cell, which reads Cariprazine hydrochloride bulk and Cariprazine capsules. A comparable entry on another agent in this chapter retains its bulk leg, so the loss was an inconsistency rather than policy, and the leg is restored. Repaired, the three instruments pass against a locator giving document, page, serial number and the pixel bounds of the indication column and row band: the

text is verbatim, it is not the primary cut short, at 215 characters against 215, and the 133-character indication cell is carried in full.

The lesson is the order of events: the defect was invisible while the cell was unreadable and appeared in the hour it became readable. An unaudited row is not a clean row, and defect counts rise when the instrument starts working.

17.8 The three records side by side, and what a complete answer contains

Three regulators, three reading dates, three indication sets.

JURISDICTION	WHAT THE RECORD HOLDS	THE SENTENCE THAT MAY BE WRITTEN	THE SENTENCE THAT MAY NOT
United States	Four label indications: schizophrenia from 13 years, acute mania or mixed episodes from 10 years, bipolar depression in adults, adjunctive therapy in major depressive disorder in adults, NDA204370, label effective 18 December 2025, read 3 August 2026	Each indication whole, with application, label date and reading date	Any wider indication, and the earliest date called a first approval
European Union	One authorised indication, schizophrenia in adult patients, from section 4.1 of the annex; authorised centrally from 13 July 2017	The indication verbatim, with the centralised-procedure scope limit	Any American indication transferred, and any claim about national authorisations this register cannot see
India	A located CDSCO entry carrying two adult indications, the bulk drug and four capsule strengths, serial number 15, dated 16 July 2021 under the heading Date of issue	The entry whole, with its serial number, document and the heading its date came from	Any negative inference from what it omits, and approval substituted for the printed heading

Reduced to five sentences:

- Partial agonism explains the shape of benefit and harm: submaximal signal at high tone, supplied signal at low tone, an activating adverse-effect profile.
- The D3 preference is a laboratory affinity ranking offered as rationale for negative and anhedonic presentations, neither outcome nor licensed claim.
- The long-lived active metabolite moves every clinical timescale outward: dose changes judged late, adverse effects resolving slowly, stopping that is not clearing.
- The indication set differs by jurisdiction and the bipolar-depression licence sits in one record, so no status sentence may be written without its jurisdiction, source and date.

- What a record omits is a fact about that record, and on the Indian corpus it cannot be converted into a statement about the country.

MNEMONIC

Preference, Persistence, Place

Preference is the D3 story, a binding ranking offered as rationale and never as licence. Persistence is the long active moiety: accumulation, delayed judgment, slow offset. Place is the jurisdiction every indication sentence must name and never leave, and why four, one and two are all correct answers here.

The receptor preference is a rationale, the metabolite makes the clock run in weeks, and the indication count is four, one or two depending on which regulator you are quoting.

18 · Lumateperone: pharmacology and indications

One capsule acts on serotonin, dopamine and glutamate at the same time, and the breadth of the licence follows from that rather than from any single receptor. Lumateperone is examinable for two unrelated reasons, and preparing one loses half the marks. The pharmacology leaves the dopamine-occupancy logic every antipsychotic algorithm is built on, which is what makes one molecule plausible across schizophrenia and bipolar depression. The regulatory record is the sharpest case in this chapter: three national records of one drug in three states, the Indian one neither an approval nor an absence. Receptor principles for the serotonin and dopamine systems belong to the Psychopharmacology I: principles and core drug classes notes and are not rebuilt here.

18.1 One molecule across three transmitter systems

The mechanistic account has four legs, taught as pharmacology rather than as anything a regulator has endorsed.

- Potent antagonism at the 5-HT2A serotonin receptor, the highest-affinity action of the molecule.
- A phase-dependent action at the dopamine D2 receptor, partial agonism presynaptically with antagonism postsynaptically, damping transmission where it is excessive rather than subtracting it everywhere.
- Inhibition of the serotonin transporter, an action shared with antidepressants rather than antipsychotics.
- Indirect enhancement of glutamatergic transmission in prefrontal circuits, dependent on the dopamine D1 receptor and expressed through NMDA and AMPA currents.

A molecule doing all four has a route into psychosis and a route into depressed mood at once, which is why the licence in 18.2 stops looking like three unrelated permissions. That is a rationale, not a demonstration: no trial evidence for this agent is held in this build, so no trial name, sample size, endpoint or effect size appears here.

The licence names none of the four legs. The current United States label calls the product **an atypical antipsychotic** and stops: no receptor, no subtype, no affinity ratio, no occupancy figure. Marks go in a predictable direction, by promoting the mechanism into the licence or by attaching an occupancy percentage to a regulatory sentence. No occupancy or affinity figure is printed here, because none is grounded in this build, and a number recalled from a review and set beside a licence is a fabricated citation whatever its provenance.

RULE **A mechanism explains why a licence might be wide; it never widens one.** Write what the molecule does, then separately what a named regulator has permitted, in that regulator's words, with the date of the reading. The two sentences must stand apart: one is pharmacology that travels everywhere, the other a permission that stops at a border.

18.2 The United States: one application, three indications, quoted whole

Exactly one jurisdiction carries a grounded status row here, and its fields are what make a claim checkable. The source is named and narrow: the openFDA interface to Drugs@FDA, one application, NDA209500, with the current structured product label effective 20 June 2025, and the status was true as of 1 August 2026. The indication text follows, with the label's pointers to its own clinical-studies sections removed and nothing else changed.

CAPLYTA is indicated for: Treatment of schizophrenia in adults. Treatment of depressive episodes associated with bipolar I or II disorder (bipolar depression) in adults, as monotherapy and as adjunctive therapy with lithium or valproate. Adjunctive therapy with antidepressants for the treatment of major depressive disorder (MDD) in adults.

Three permissions, and they are not parallel.

LICENSED INDICATION, UNITED STATES	MAY BE GIVEN ALONE	MAY BE GIVEN WITH
Treatment of schizophrenia in adults	No partner drug named either way	No partner drug named either way
Depressive episodes associated with bipolar I or II disorder (bipolar depression) in adults	Yes, stated as **monotherapy**	Also stated as **adjunctive therapy** with lithium or valproate
Major depressive disorder in adults	No such wording appears	Adjunctive therapy with antidepressants only

The asymmetry is the mark. Bipolar depression is licensed both ways; major depressive disorder one way only, as an addition to an antidepressant. The bipolar row covers bipolar I and bipolar II together, which much of the literature a candidate has read does not. Note what the schizophrenia row does not say: an adult indication, no adolescent extension, no line of therapy, no comparator, no place in any algorithm.

Then the field most often misreported. The earliest approval-status date for this application is **20 December 2019**, and the field is named **the earliest approval-status date in the record**. Whether that is a first marketing approval, a supplement or something else, the record read here does not say. This ledger has already retracted one drug's date after publishing it under a field

name that promised more than the field held. Print a date with the name of the field it came from.

Two disciplines travel with the American row. The interface is an agency-operated mirror of Drugs@FDA and of the label archive rather than the Drugs@FDA application itself, and the indication wording is current label text as filed, which is not a statement about any other jurisdiction and must never be carried across one. And the row was tested rather than trusted. Three text instruments ran against the primary document, checking that the row text is verbatim, that it is not the primary text cut short and closed so that nothing signals the loss, and that the counts of indication statements match. All three passed, at 848 characters against 848, and five indication statements against five. Five rather than three counts the label's highlights repeat of the same list.

GOOD TO REMEMBER

Licensing bipolar depression both as monotherapy and as an addition to lithium or valproate answers the question the ward asks, which is whether the existing mood stabiliser has to come off. On this wording it does not. What the licence does not answer is whether adding this drug beats optimising what is there, because no comparative, sequencing or head-to-head result for this agent is held in this build. Answer the licensing question from the label; refuse the comparative one aloud.

TRAP **Carrying the monotherapy permission across from bipolar depression to major depressive disorder.** The wording gives bipolar I and bipolar II depression as monotherapy and as adjunct; it gives major depressive disorder as adjunctive therapy with antidepressants and in no other form. Writing that the drug is licensed for depression, unqualified, merges two indications with different permissions and different partner drugs, and usually exports an American permission to whichever country the examiner had in mind.

18.3 The European Union: an absence in a register that records refusals too

A declared gap is a statement about a search and not about a product, so its strength is its instrument's strength. The instrument is the European Medicines Agency's own published medicines register, generated by the agency on 1 August 2026, downloaded whole and scanned across all 39 columns of all 2730 medicine rows, where this active substance returns zero rows under any of its names.

What lifts that above a shrug is the range of outcomes the register carries. It records Authorised, Refused, Application withdrawn, Withdrawn, Lapsed, Expired, Revoked, Suspended and Opinion outcomes, so the absence means no central application of any kind was ever concluded, including one refused or abandoned partway. The missing document is an agency public assessment report for the substance or a European Commission implementing decision, and the phrase approved internationally is not a permitted substitute for a named regulator.

Now the limit, where an honest answer separates itself from a confident one. The register covers the European Union **centralised procedure** only. A product authorised nationally in one or more member states through the decentralised or mutual-recognition procedure would never

appear in it, so absence here means no centralised procedure of any kind is on record, which is narrower and weaker than not available anywhere in the Union. That limit sits on every European gap row and does not bite equally on all; on this agent it is graded as biting moderately. Two weaknesses travel with the finding. The confirming second instrument ran on all eight grounded European rows and on none of the eight gap rows, so no European gap here has independent corroboration, and the assessment-report page is in any case a second rendering of the same source, because the register is generated from those pages.

The defensible European sentence is short and worth writing whole: no European Union centralised procedure of any kind is on record for lumateperone as of 1 August 2026, which does not exclude a national authorisation this instrument cannot see.

18.4 India: a committee that recites an approval no list prints

The Indian entry is the one an Indian examination cares about, and it is the only row in this chapter that is neither a grounded approval nor a declared gap. Its state is partial, on a supportive source only. One affirmative document exists, a committee minute rather than a permission. Minutes of the CDSCO **subject expert committee** for neurology and psychiatry, dated 16 April 2025, record verbatim that the committee noted that drug Lumateperone Capsules 42mg is already approved in the country on 26.12.2024. The date recited is **26 December 2024**, and the strength named, **42 mg**, is part of the product designation in that minute and not a dosing instruction; no dosage text for this agent is held in this build.

What the row does not hold is a **marketing authorisation**. Read the next sentence twice, because collapsing it is the most expensive error on this item. Subject expert committee minutes are a supportive source only, so the row licenses the statement that CDSCO committee minutes of 16 April 2025 recite such an approval, and it does not license the statement that CDSCO approved lumateperone on 26 December 2024. Two claims about different things, one about a document and one about a regulatory act, and only the first is in evidence. The second would rest on the Drugs Controller General of India, whose permission was searched for and not located.

The primary record was sought in three forms and found in none. The approved indication wording was not located at all, so the row records the fact of an approval being recited and not its indication. No CDSCO permission and no approval-list entry appears in the 66 published approval-list documents. A run-time scan of all 121 extracted CDSCO texts found the drug's name in none of them.

The list that should have carried it demonstrably covers the date. The list declaring calendar-2024 coverage carries entries dated 17 December 2024 and 30 December 2024, and declares its own range in its running header as covering 1 January 2024 to 31 December 2024, so it is not truncated before 26 December. Two corrections to that argument sit in the row and are worth learning as technique. An earlier version cited entries dated 17, 30 and 31 December 2024, but 31 December occurs only inside the page header and is not an entry; the declared coverage range is the stronger evidence anyway, because it excludes truncation by the document's own statement rather than by inference from its last row. A second leg, that the list carries roughly 15 numbered

rows, was withdrawn as false: there is no single calendar-2024 list, the corpus holds five snapshot vintages across two CDSCO series, and the one declaring full-year coverage carries 30 recoverable rows to a highest serial of 31.

IN INDIA

Write it in this order and stop. CDSCO subject expert committee minutes of 16 April 2025 recite an approval of Lumateperone Capsules 42 mg in the country on 26 December 2024; committee minutes are a supportive source and not a marketing authorisation; the permission, an approval-list entry and an approved label were each searched for and none was located, so the approval is recited here and not established here. Do not write that CDSCO approved the drug on that date, which upgrades a recital into a ruling. Do not write the categorical negative either, because no instrument in this build supports it. For a patient, availability is checked against a current label or a pharmacy record, never inferred from a committee minute.

18.5 Why this one row licenses every declared gap in the India column

This row does more work than its own drug. It is the strongest of the three legs establishing that the CDSCO published approval-list corpus is incomplete, because here CDSCO contradicts itself: its own committee recites an approval that appears in none of the 66 published lists, and the list that should have carried it declares coverage of the date. Every other Indian entry in this chapter depends on that incompleteness, so one compact agent profile is why the rest of the column may be written as declared gaps at all.

The second leg is an era mismatch, and it bears on a chapter of recent approvals: every drug this corpus is known to find sits in a document dated 2023 or earlier, and independent sensitivity in the 2024 to 2026 window is 0 of 1, that one being this drug, whose Indian approval is established from outside the corpus and which the corpus misses. The ledger refuses to count the one agent here that the corpus does find as a second data point, because scoring a corpus hit as proof that the corpus finds hits is circular, and the positive-control tally of 17 of 20 may not warrant a gap on any agent approved after 2023.

848 of 848	0 of 2730	0 of 66	0 of 121
CHARACTERS, US ROW VERSUS LABEL	EU REGISTER MEDICINE ROWS MATCHING	CDSCO APPROVAL LISTS WITH AN ENTRY	EXTRACTED CDSCO TEXTS NAMING THE DRUG

Quote the conclusion rather than improving on it. Absence from this corpus is evidence and not proof; no India row in this build asserts that a drug is not approved in India; each asserts instead that no CDSCO approval record was located, and names what was searched. A categorical negative is a claim about a country, and the instrument only ever supported a claim about a search.

18.6 A date and the label on the date are one record

One correction inside this row is the most transferable thing on the item, and it was made by hand after independent adjudication. The field describing what kind of date 26 December 2024 is once carried the wording the grounded Indian rows carry, which means an entry in an approval list. Anything joining the date to that description would have emitted a false Indian approval

sentence, correctly formatted and entirely wrong. It was raised as a blocker twice by independent cross-family review, and the row's state was left unchanged, because partial on a supportive source was already correct. A date and the label describing it are one record and must be corrected as one unit, and a string true in one row is not therefore true in another.

The falsification pass makes the point mechanically, and its verdict on the Indian row is a term worth owning: **unresolved extraction**. All three text instruments returned it rather than a pass, because this row cites no approval-list entry and therefore offers no indication cell to read. The finding was recorded as measured rather than assumed, through the run-time scan of all 121 extracted texts and a second source check across five cached committee-minute files, in one of which the quoted sentence appears verbatim. The discipline is written into that file in one line: a failed extraction is unresolved extraction, never a pass. An instrument that could not read is not an instrument that found nothing.

MNEMONIC

Recited, not Ruled

Recited is what a committee minute does: it reports an approval. Ruled is what a permission, an approval-list entry or an approved label would do: it constitutes one. The Indian row holds the first and not the second, so the sentence it licenses names the minute, its date and its verb. Naming CDSCO as the actor on 26 December 2024 upgrades a recital into a ruling.

Reduced to three sentences, the item is this.

- The mechanism acts on serotonin, dopamine and glutamate at once, and is a rationale for the licence, never a substitute for reading it.
- The American record is grounded and quotable: schizophrenia in adults; bipolar I or II depression alone or added to lithium or valproate; major depressive disorder only as an addition to an antidepressant.
- The Indian record is a recited approval and not an established one, and that single row is the evidence that the published approval-list corpus is incomplete.

One molecule across three transmitter systems, one licence quoted whole and left inside its own country, and one Indian record that recites an approval no published list prints.

19 · Zuranolone as an individual agent profile (cross-reference to recent advances)

This is the prescribing half of zuranolone, and it is short on purpose. The mechanism, the class lineage, the postpartum against major depressive reading and the regulatory chronology belong to the zuranolone and neurosteroid antidepressants item in this chapter. What follows is the agent-reference residue: the dose, the pharmacokinetics, the interactions, and the points where two labels describe one molecule differently.

19.1 Identity, and the border running under every line below

Zuranolone is an oral neuroactive steroid, classified in the European document under psychoanaleptics, other antidepressants, with ATC code N06AX31, supplied as 20 mg, 25 mg and 30 mg hard capsules. Two regulatory records are held here and a third is empty. The United States label reads that ZURZUVAE is indicated for the treatment of postpartum depression (PPD) in adults, on Drugs@FDA application NDA217369, whose earliest approval-status date in the record is 4 August 2023, status true on 1 August 2026. In the European Union the medicine is Zurzuvae, EMEA/H/C/006488, holder Biogen Netherlands B.V., authorised throughout the Union from 17 September 2025; its indications section names one indication only, quoted verbatim in the neurosteroid item and deliberately not paraphrased here. For India, no CDSCO marketing approval record was located as of August 2026, on the authority of the Central Drugs Standard Control Organisation and the Drugs Controller General of India.

Everything below is label content, and a label describes what one regulator permitted inside its own territory. A dosing table is not a permission, and no figure here crosses a border.

19.2 One dose, one evening, fourteen days, and what moves it

SITUATION	DOSE	WHAT THE LABELS SAY
Standard	50 mg orally once daily in the evening for 14 days	The European document calls this a single course; the United States label states that safety and effectiveness beyond 14 days in a single treatment course have not been established
Not tolerated	40 mg once daily, same 14 days	European trigger is failure to tolerate 50 mg, United States trigger is depressant effects arising inside the period. The European document adds that dosing may be stopped without down-titration
Strong CYP3A inhibitor	**30 mg once daily through the 14 day treatment period**	Both. The United States label adds that no dosage modification is recommended with a moderate CYP3A4 inhibitor
Moderate or severe renal impairment	30 mg, same period	European: moderate eGFR 30 to 59 mL/min, severe under 30 not requiring dialysis. United States: moderate or severe, eGFR under 60 mL/min/1.73 m2. Mild, 60 to 89, unchanged in both
Severe hepatic impairment	30 mg, same period	**Child-Pugh class C** only; class A and class B unchanged in both
Missed evening dose	Nothing extra	Take the next dose at the regular evening time, take no additional capsules the same day, and continue to the end of the course. Both

The capsule arithmetic goes wrong on the prescription, not in the answer. Fifty milligrams is two 25 mg capsules, forty is two 20 mg capsules, thirty is a single 30 mg capsule, so every dose change is a change of strength, not of count.

The administration instruction is an exposure requirement. Capsules are swallowed whole in the evening with fat-containing food, as either a meal or a snack, quantified by the United States label as 400 to 1,000 calories and 25 to 50 per cent fat and left by the European document as everyday examples. Zuranolone may be given alone or alongside stable background oral antidepressant therapy, so it is neither an add-on by licence nor a replacement for one.

RULE **Every dose in this entry is a fourteen-day dose.** There is no titration upward, no loading and no maintenance limb. Both labels specify each reduced dose for the whole treatment period rather than for one evening, so a patient started at 30 mg for an inhibitor, a kidney or a Child-Pugh class C liver stays there to the end. The dose is decided once, at the prescription; the only trigger arising later is tolerability.

19.3 The pharmacokinetics that explain the instructions

Two of these numbers are why the patient is told anything at all. Each is printed with the document it came from, because the two labels do not always use the same measure.

PARAMETER	VALUE	DOCUMENT, AND ANY DIVERGENCE
Time to peak concentration	**5 to 6 hours after the dose**	Both
Steady state	Reached in 3 to 5 days, accumulation about 1.5-fold	Both
Distribution	Volume of distribution greater than 500 L; plasma protein binding greater than 99.5 per cent	Both; the European document adds that the volume is dose-independent
Terminal half-life	**Approximately 19.7 to 24.6 hours in an adult population**	Both
Mean apparent clearance	32.7 L/h European, 33 L/h United States	One quantity, two roundings; quote the one whose document you name
Excretion	45 per cent in urine and 41 per cent in faeces, both as metabolites, with negligible unchanged drug in urine and under 2 per cent in faeces	Both. Renal impairment matters through exposure, not unchanged excretion
Food effect at 30 mg	European: peak up 228 per cent and AUC up 55 per cent low-fat, 334 and 90 per cent high-fat, against fasting. United States: peak up about 3.5-fold and AUClast about 1.8-fold low-fat, 4.3-fold and 2-fold high-fat	Two documents, two metrics; time to peak unaffected in both

50 mg	30 mg	5 to 6 h	19.7 to 24.6 h
THE ONLY DAILY DOSE	THE REDUCED DOSE, THREE TRIGGERS	TIME TO PEAK CONCENTRATION	TERMINAL HALF-LIFE

A half-life near a day with once-daily dosing keeps the course at steady state throughout, so the twelve-hour restriction both labels attach to each dose is not a wash-out and does not lengthen as the fortnight goes on. It is set against a performance threshold, worked through in the neurosteroid item, not against elimination.

The exposure figures behind those thresholds: in Child-Pugh class C the peak concentration was 24 per cent lower and AUCinf 56 per cent higher than in matched healthy subjects, while class A and class B were unchanged, which is why only the severe class moves the dose. Exposure rose in moderate and severe renal impairment, and the drug has not been studied at an eGFR under 15 or in patients requiring dialysis. No adjustment is made for weight, race or age; Black or African American subjects showed a 14 per cent higher apparent clearance, which the European document itself calls not clinically meaningful. There is no geriatric experience in this indication.

19.4 Interactions: a victim of CYP3A, almost never a perpetrator

Metabolism is extensive with CYP3A primary, so the interactions worth knowing run one way. A **strong CYP3A inhibitor** raises exposure, quantified as AUCinf up 62 per cent with itraconazole, and the answer is 30 mg rather than avoidance; the European examples are protease inhibitors, azole antifungals and some macrolides, with grapefruit avoided for the same reason. Induction reverses both magnitude and management: AUCinf falls 85 per cent with rifampin, and inducers are avoided outright with no dose-adjusted alternative. The European list names carbamazepine, phenobarbitol, phenytoin, primidone, rifampicin, St John's Wort and efavirenz; the last three are the ones candidates omit.

In the other direction the drug is quiet. Repeated zuranolone did not alter exposure to simvastatin, a CYP3A substrate, or bupropion, a CYP2B6 substrate, and it is not a substrate of P-glycoprotein, nor expected to inhibit the common cytochromes or the major transporters at clinically relevant concentrations. The one important interaction that is not pharmacokinetic at all is with other central nervous system depressants: repeated 50 mg daily doses with alcohol or alprazolam produced increased impairment of psychomotor performance, yet the United States pharmacology records that alprazolam and ethanol did not alter zuranolone exposure, and zuranolone did not alter theirs. The additive effect is pharmacodynamic, so no concentration would predict it, and the management is dose reduction where the combination is unavoidable.

19.5 Where the two labels do not say the same thing

Contraindications first, because the divergence is total. The European document contraindicates hypersensitivity to the active substance or excipients, and pregnancy. The United States contraindications section reads, in its entirety, None. That is not a disagreement about foetal risk: both advise effective contraception during treatment and for one week after the final dose, and the United States label carries an embryo-fetal toxicity warning on the same animal findings. One regulator used a contraindication, the other a warning.

Dependence next. The European special-warnings section states that, on clinical-trial data, zuranolone has low physical dependence potential. The United States label states that it may produce physical dependence, and prints the evidence: in healthy subjects given 50 mg for 6 to 7 days, discontinuation was followed by insomnia, palpitations, decreased appetite, nightmare, nausea, hyperhidrosis and paranoia, all mild to moderate. In that jurisdiction zuranolone is a **Schedule IV controlled substance**, a classification belonging to United States law that must not be carried elsewhere. Quoting either sentence and calling it the position on **physical dependence** quotes half a field.

The United States label also carries the antidepressant-class warning on suicidal thoughts and behaviour, with pooled figures showing 14 additional patients per 1,000 under 18 and 5 additional at 18 to 24 against 1 fewer at 25 to 64, adding that zuranolone does not directly affect monoaminergic systems. The European special-warnings section, read whole, carries five headings, on driving, depressant effects, interactions, abuse and dependence, and excipients; suicidality is not among them.

The paediatric sentences differ in kind and must not be merged. The European document records that safety and efficacy in postpubertal females less than 18 years old have not been established, and that there is no relevant use in prepubertal females. The United States label records only that safety and effectiveness in paediatric patients have not been established. Reading either as a gloss on the other manufactures a floor neither document sets.

TRAP **Quoting a frequency without its population.** The European pooled figure for somnolence among subjects given 50 mg is 26.5 per cent; the United States table for the same dose in its Study 1 gives 36 per cent against 6 per cent on placebo. Both are correct and both are 50 mg figures; the poolings differ. The trap runs the other way on tolerability: the European frequency for dose reduction or interruption is 14.3 per cent, while the United States figure for reduction alone is 16 per cent, so the narrower construct carries the larger number, which is possible only because the denominators are different patients.

19.6 The prescribing points, in the order they are decided

One intentional overdose was reported before marketing: 330 mg, about 6.5 times the maximum recommended dose, producing an altered state of consciousness that resolved the next day after intravenous fluids. There is no specific antidote and management is supportive.

- Decide the dose before the first capsule: a strong CYP3A inhibitor, a moderate or severe eGFR, or a Child-Pugh class C liver each fixes the course at 30 mg, and all three are readable from the chart.
- Check inducers separately, because they are an avoidance and not a dose problem: carbamazepine, phenytoin, phenobarbitol, primidone, rifampicin, St John's Wort, efavirenz.
- Counsel the evening, the fat and the twelve hours as one instruction with three parts.
- Name the jurisdiction whenever a status, a schedule, a contraindication or a frequency is quoted, because each of those four differs between the documents or exists in only one.

GOOD TO REMEMBER

Three of the four reasons to reduce are known before the prescription is written and all three land on the same 30 mg capsule; the fourth, tolerability, lands on 40 mg and arises during the course. Settle renal function, liver status and the interacting list at prescribing, and keep the 40 mg step in reserve for the first two evenings.

IN INDIA

Every figure in this entry comes from two foreign labels, and a label describes a permission granted elsewhere. As of August 2026 no CDSCO marketing approval record was located for zuranolone; the missing document is an affirmative CDSCO or DCGI new-drug permission, an approved Indian label, or an entry in an official approval list. A CTRI registration, expert-committee minutes and a brand in a distributor catalogue are none of them an Indian marketing authorisation. This item holds no instrument that measured Indian availability in either direction, and the measured limits of the approval-list corpus are set out in the neurosteroid item.

MNEMONIC

Evening, Fat, Fourteen, Thirty

Evening is when it is given and why the restriction runs into the next morning. Fat is the condition of absorption. Fourteen is the whole exposure, with no titration and no second course. Thirty is the reduced dose and its three triggers, all knowable before the first capsule.

One dose for fourteen evenings, one reduced dose for three chart-readable reasons, and a jurisdiction attached to every status, schedule, contraindication and frequency you quote.

Consolidate and apply

20 · Worked vignettes

A vignette in this chapter is almost never a question about a drug; it is a question about a record. The stems that carry marks open with somebody wanting something: a patient who has read about a new treatment, a colleague who prescribed it in another country, a committee asked to stock it. Then the stem plants the fact that decides the answer, and in this territory that fact is almost always a jurisdiction, a document, a date, or a kind of permission. What the closing vignettes rehearse is the reading skill the chapter was built to install: finding the status question inside a clinical stem, and answering it at the grain the evidence actually supports.

The discipline that earns marks is refusing to answer a regulatory question with a pharmacological one, and refusing to answer it about the wrong country. A patient asking whether they can have a treatment they have read about is not asking what it does at the receptor; they are asking whether a named regulator has authorised it for them, here, now. A candidate who supplies the American record for a patient in an Indian outpatient clinic has answered the right question in the wrong place, which costs more because it looks like knowledge. Where a stem turns on mechanism, or on choosing within a class, that reasoning belongs to the Mood disorders: pharmacotherapy notes and the Psychopharmacology I: principles and core drug classes notes. This chapter supplies the layer above the prescription: whether the thing may be prescribed at all, on whose authority, and how to say so when the answer is not known.

20.1 Four questions locate the status claim

Before deciding what a recent-advances stem is asking, run the same four questions across it. Between them they cover almost every way an examiner can turn a request for a new treatment into a regulatory problem, and each opens onto a different part of this chapter.

Border	Document	Date	Verb
WHICH JURISDICTION IS ASKED	WHAT WOULD SETTLE IT	WHEN IT WAS TRUE	WHAT KIND OF PERMISSION

Border asks which country's record the stem is asking about, and that is set by where the patient is, not by where the literature was written. Document asks what record would settle the question and which body issues it: a published approval list, a current approved label, a public assessment report, a commission decision. Date asks when the status was last checked, because a regulatory statement with no date is a claim about the present made from an unspecified past. Verb asks what kind of permission is described, since permission to run a trial, an approval-list entry, a recital inside committee minutes and a product being supplied are four different facts about one agent. A well-built stem makes one of the four decisive and leaves the others as scenery.

RULE **A status sentence with no jurisdiction, no source and no date is not an answer.** Those three fields, with the indication itself, are what a regulatory claim here has to carry, and a sentence missing any of them cannot be checked. "It is approved" is not a shorter version of the answer; it is a different and unsupportable claim, because it silently asserts the same status everywhere. If you cannot give all three, say what was searched and what would settle it: a complete answer of a different shape, not a weaker version of the same one.

20.2 The vignette that resolves on a border

The commonest stem names an agent and asks whether the patient in front of you can have it, and the candidate answers from whichever country's record they happened to revise. The reading move is to fix the border first, and to treat the answer as a property of the pair, agent and jurisdiction, rather than of the agent. Regulatory maturity does not run in one fixed order, and this chapter holds the counter-example. One of the orexin receptor antagonists holds an Indian approval-list entry with an insomnia indication, dated **1 July 2021** under the printed column heading Date of issue, while for the same substance no European centralised procedure of any kind was found on record, both status true as of 1 August 2026. A candidate who reasons that an agent must be available in India because it is available in the United States, or that it cannot be because Europe has no record of it, has assumed a ladder that does not exist.

The other half of the move is knowing what a European absence is worth. The instrument behind it is the agency's own register of centrally authorised medicines, which records refusals, withdrawals, lapses, revocations and suspensions alongside authorisations, so an agent absent from it has never concluded a central application of any kind. That is a real finding, and narrower than it sounds, because the register covers the centralised procedure only and a product authorised nationally through the decentralised or mutual-recognition route would never appear in it. One further weakness belongs in a careful answer: the confirming second instrument ran on

all eight grounded rows and none of the eight absences, so no European absence here has independent corroboration.

20.3 The vignette in which the record was searched and not found

A second family hands the candidate an agent for which nothing was found, and the examinable content is entirely how that is said. A patient asks for an intranasal treatment they have read about; the Indian record was searched, and no approval was located. The wrong answer converts an absence of record into a regulator's negative finding. The right answer is a **declared gap**: name what was searched, name the authority that would settle it, and name the document that is missing. For the Indian column that means three instruments, namely a corpus of published approval lists, a search of the regulator's own domain in which every returned document is read rather than counted, and a search of the national trial registry run with a positive control. **Marketing authorisation** is what would settle it, and nothing softer will do.

The instrument's blindness has been measured, which is why this is not pedantry. The published approval lists are demonstrably incomplete, shown most strongly by the regulator contradicting itself: its own expert committee minutes of April 2025 recite an approval for an agent that appears in **0 of 66** of those documents, and the list declaring calendar-2024 coverage states its own range in its running header, so it is not truncated before the recited date. A corpus that misses a real approval cannot then be offered as evidence that some other approval does not exist. Absence here is evidence and not proof, and the sentence to write is that no approval record was located, naming what was searched.

IN INDIA

Write it in this order. The authority is the Central Drugs Standard Control Organisation and the Drugs Controller General of India. Expert committee minutes, **clinical-trial no-objection certificates**, International Cell written confirmations, CTRI registration, a brand listing, pharmacy availability and manufacturer copy are none of them marketing authorisation. So the honest sentence to a patient or a referring colleague is that no Indian approval record was located as of the date you checked, that what would close the question is a CDSCO or DCGI new-drug permission, an approval-list entry or a current approved Indian label, and that a pharmacy being able to source something is a different fact.

TRAP **Converting a search that found nothing into a regulator's negative finding.** Writing that an agent is not approved in India, or not approved in the United States, asserts that a body has stated something no instrument here asked it. The probes measured named routes and named documents; they are silent on what a regulator would say if asked directly. Worse, an American gap in this build has already been wrong once: one agent's record carries no cross-reference block at all, so three of five routes returned not-found on it while the active-ingredient route returned the record. A declared absence is written at the strength the instrument supports and never one notch harder, and an examiner marks the notch.

20.4 The vignette in which the approval stands and the product does not

A stem describes a woman with severe postpartum depression whose team has read about an intravenous neurosteroid and wants to know whether it can be arranged. An approval and a marketing status are two separate facts about one product. The American record for the first neurosteroid in this class carries an approval dated 17 June 2019 together with marketing status Discontinued. Both fields, approval and Discontinued marketing status, sit on that one record, which carries nothing relating them, so an answer gives both and reasons from neither to the other. Naming the approval and dropping the Discontinued marketing status tells the family half of what the record says. Saying the drug was withdrawn as unsafe is the more instructive overreach, because it reads a reason into a field that gives none: the Discontinued marketing status is a fact about the product rather than about the approval, and no field on this row states why marketing stopped or names any action taken about it by a regulator or a sponsor, so a candidate who supplies one has left the record, in either direction. Say that the approval stands on the record and the marketing status is Discontinued, and stop where the record stops.

The same row, the Discontinued one, carries a second object worth a sentence in a viva. Its date of approval and the date its approval letter was signed are different quantities, because the approval letter sets the date of approval as the date of a later scheduling notice, and the earlier date has been retracted from the field a reader would take a first-approval date from. Publishing a signature date under a first-approval field name is a true date of one thing wearing the field name of another, which is the commonest way a correct number becomes a false claim. Where an answer usually lands instead is the oral agent in the same class, whose rows sit in this chapter's neurosteroid section.

20.5 The vignette that resolves on the kind of document

Four different things get called approval in conversation and only one is a marketing authorisation, so a stem can be built entirely out of the difference. The first is a trial permission. The Indian file for the muscarinic combination is documented at pre-approval stage: no-objection certificates name the combination, and committee minutes for neurology and psychiatry dated 30 May 2025, 10 December 2025 and 22 December 2025 record a named sponsor presenting Phase III protocols, with the committee recommending permission to conduct the trial. That is affirmative evidence about the file and it is not an approval. The second is a recital: minutes of a **subject expert committee** mentioning an approval for another agent license the statement that those minutes recite such an approval, and do not license the statement that the regulator approved the drug on that date. The third is supply, a fact about a market. The fourth is a strength printed beside a product name in a list, which is part of the product's identity and not a dosing instruction.

The value sits in the three errors this catches. A candidate who writes that the muscarinic combination is approved in India has upgraded a trial permission into a marketing authorisation. One who writes that the other agent was approved on the recited date has upgraded a recital into the permission itself, when the primary permission or label was never located. And one who writes that either agent is unapproved has downgraded affirmative documentary evidence into a

negative nobody established. A committee recommending that permission be granted is a step towards permission and not the permission, so say which document you are describing before you say what it shows.

20.6 The vignette that resolves on a record that is authoritative and wrong

The next family teaches a reading habit rather than a fact. A record can come from the right regulator, cite the right document and still say something false, and both shapes it takes are silent. The first is truncation. The European indication for the oxybate medicine reads Treatment of narcolepsy with cataplexy in adult patients, adolescents and children from the age of 7 years. An earlier stored version of the same field read "Treatment of narcolepsy with cataplexy in adult patients." The words "adolescents and children from the age of 7 years" had been dropped and the surviving prefix closed with a full stop the source does not have. Nothing looked wrong, because a cut closing with valid punctuation is invisible to the eye. The vignette that follows is paediatric: a child with narcolepsy and cataplexy, and a candidate whose remembered indication has quietly excluded them.

The second shape is a wrong product wearing the right citation. The cited application and the date were both correct and both belonged to one product; the indication text was not, because it was the text of a different application by a different sponsor, whose label reads almost identically in structure. A reader comparing structure or length would see nothing. Both defects were found the same way, by somebody reading a stored field, noticing it did not close properly or did not name the product it claimed to, and refusing to complete it from memory. That refusal is the transferable skill: check the product's identity against the application's, quote an indication whole or say you are paraphrasing, and treat a quotation ending on a tidy full stop as unverified.

20.7 The vignette that resolves on what the evidence does not license

Not every stem here is regulatory. The other family gives a finding and asks what follows, and the mark sits in the inferences refused rather than the ones drawn. The umbrella review of the serotonin literature is the standing example. What it examined and what it reported are one question; what it licenses is another; and it does not license the conclusion that antidepressants do not work, that depression is not biological, that serotonin has no role in mood, or that evidence about aetiology settles a question about efficacy. Aetiology and efficacy are separate claims answered by separate designs, and a stem offering the review as grounds for stopping a patient's treatment is testing exactly that separation.

The same discipline runs through the newer research territory. A psychedelic-assisted trial is an intervention package, drug with preparation, setting and psychological support, so a question about what worked is a question about which component the design can separate, with expectancy, functional unblinding, therapist effects and the exclusion of higher-risk patients all on the appraisal path; and an access arrangement for a named patient is not a product approval. A stem describing an algorithm, a candidate biomarker or a prescription application is asking which rung the evidence has reached, across analytical validity, clinical validity, clinical utility, regulatory authorisation and routine deployment. No rung is inherited from the one below, and a **measured absence** of evidence at one rung is not a claim about the rung beneath. Those

arguments are developed in this chapter's serotonin-hypothesis, psychedelic-therapy, machine-learning, biomarker and digital-therapeutics sections; an answer here names the rung and stops.

20.8 Look-alike stems and the answering spine

Consolidation is easiest in pairs. The table sets out the stems most often confused here, against the fact that separates them and the axis the reasoning lands on.

TWO STEMS THAT LOOK ALIKE	THE FACT THAT SEPARATES THEM	THE AXIS IT LANDS ON
An agent available in one country, against the same agent asked about in another	Which jurisdiction's document was actually read, since the answer belongs to the pair and not to the drug	Border
An absence of record, against a regulator's negative finding	Whether any body has stated that nothing exists, or only that named searches found nothing	Document
An indication quoted whole, against one truncated at a full stop the source does not have	Whether the population at the end of the sentence survived the quotation	Document
A clinical-trial permission, against a marketing approval	What the document permits: conducting a study, or marketing for an indication	Verb
Committee minutes reciting an approval, against the approval record itself	Whether the permission, the list entry or the approved label was ever located	Verb
A status checked this month, against one carried in from memory	Whether the claim carries the date on which it was true	Date

With the axis identified, an answer that earns marks follows four moves in order.

- Fix the border. Name the jurisdiction the question is being asked in, before naming any status at all.
- Name the document that would settle it, and the body that issues that document.
- Give the verb its correct strength: authorised, permitted for a trial, recited in minutes, supplied, discontinued.
- Say what the record does not settle, and what you would have to read to close it.

That fourth move is where regulatory reasoning becomes visible to an examiner. A candidate who answers that a drug is approved has made a claim about a molecule; one who answers that a named regulator authorised a named product for a named indication, on a document they can identify, true as of a date they can give, and that the Indian record was searched on three named instruments with no approval located, has made a claim about a record, and a claim about a record can be checked. Every vignette in this territory rewards the second answer.

Answer at the grain of agent and jurisdiction, with the document, the date and the strength of the verb attached, and turn a search that found nothing into a declared gap rather than into a negative.

21 · Consolidation and quick review

This page teaches nothing and asks everything. The eighteen sections behind it have each been taught once, in the section that owns them, and consolidating them is not rereading them. It is finding out what can be produced from a blank page with the book shut, because that is the only form in which any of this will be available in a viva. What follows is a list of prompts rather than a summary: each names a section and the answer it should be possible to build cold.

Two habits make it easy to consolidate badly. One is revising the chapter as eighteen unrelated topics, which yields eighteen lists and no method. The other is holding a regulatory position as a property of a molecule; recalled that way it is unmarkable rather than merely weak, and it is usually wrong about which regulator it belongs to.

21.1 The two things to hold before any prompt below

The first is the unit of recall. Nothing here is held at the level of an agent: a status is held at the level of the **agent and jurisdiction cell**, and the cells of one agent are independently sourced and free to disagree. The ledger behind these notes carries **48** such cells, and this chapter holds pairs in which one cell is grounded while its neighbour is a declared gap. A single verdict about an agent is therefore not a compressed answer but a claim about cells that were never read. Producing one is the **border crossing** this chapter exists to prevent, and it costs marks both ways: an approval reported for a regulator holding no record, and a categorical negative written where only an instrument was silent.

The second is the audit. Four questions to run against any status answer the moment it is written, one per field such a sentence carries, the last being the date the cell was true, which these notes record as status true as of.

- Which regulator is this sentence about, and does the sentence itself say so?
- Where did the wording come from, and is it the document's own wording rather than a paraphrase that has quietly tightened it?
- On what date was this true, and does that date appear beside the claim?
- If the answer is a negative, what was actually searched, and whose authority would the claim finally rest on?

RULE **Recall by cell, never by agent, and produce the sentence rather than the verdict.** An answer naming an agent and a status has answered a question nobody can mark. An answer naming the regulator, the wording, the document and the date has answered the question asked, even when what it reports is a gap.

18	48	1	0
SECTIONS TO PRODUCE COLD	AGENT AND JURISDICTION CELLS BEHIND THEM	CELL A STATUS SENTENCE MAY SPEAK FOR	EVIDENCE LANES THIS CHAPTER HAS RUN

21.2 Twelve prompts for the recent advances band

One prompt per section. The middle column is what to produce; the right-hand column is what an otherwise fluent answer still misses, which is where the marks sit. No cell supplies the answer it asks for, with two exceptions the chapter's own rules force: a declared gap can be named only by naming what was searched, and one product's approval can be named only beside its marketing status, which a hard rule of this chapter requires in the same sentence as the product.

SECTION	PRODUCE FROM A BLANK PAGE	STILL INCOMPLETE WITHOUT
Esketamine and intranasal therapy	The indication each regulator's own record carries in its own words, the monotherapy and adjunctive positions separated per regulator, and any limitations of use the United States wording attaches	A jurisdiction on every clause, and the Indian cell written in whichever of the three states its own record supports
Zuranolone and the neurosteroid antidepressants	The neurosteroid rationale, the postpartum and the major-depression evidence read separately, and what brexanolone teaches about an approval that stands while the product's United States marketing status is recorded as Discontinued	The two populations kept apart, and the reason a date of approval can be set by something other than the approval letter
Rapid-acting antidepressants as a class	What makes the category a category rather than a list, and its members placed against onset, durability, evidence maturity, treatment setting and regulatory maturity	Each axis ordered qualitatively with its basis named, and not one figure on any of the five
Dual orexin receptor antagonists and sodium oxybate	The agents cell by cell, and the pair of cells that defeats reading the three regulators as a ladder of maturity	Whichever population a quotation of the European oxybate wording silently drops when it is cut at a full stop
Xanomeline, trospium and the muscarinic antipsychotics	Why a peripheral antagonist is paired with the agonist, and where the strategy sits beside rather than inside the dopamine algorithm	Which of the two named substances the Indian approval-list corpus cannot display, and what that does to any negative about the pair
Endoxifen and the protein-kinase-C-targeted agents	The Indian cell and the United States cell as two sentences that never meet, with protein kinase C inhibition named as a rationale and not a result	The refusal to merge the two cells into one agent-level statement; the United States cell given as the named routes that were run against Drugs@FDA and found nothing, never as a verdict, and any negative left to rest on the Food and Drug Administration's own reading of that register; and no comparison against an established mood stabiliser
Anti-amyloid monoclonal antibodies	The class of claim these agents exemplify, why this chapter prints no position for them, and what an Indian old-age service would have to establish first	The delegation named as one, the material handed to the Dementias and neurocognitive disorders notes, and no status and no date anywhere

SECTION	PRODUCE FROM A BLANK PAGE	STILL INCOMPLETE WITHOUT
The serotonin hypothesis debate	What the umbrella review examined and what it reported, then the four inferences it does not license	All four stated, and evidence about aetiology kept apart from evidence about efficacy
Psychedelic-assisted therapy	The intervention as a package of drug, preparation, setting and psychological support, with psilocybin and MDMA both named and neither given a status	Expectancy, functional unblinding, therapist effects and the exclusion of high-risk patients marked on the appraisal, and access arrangements kept apart from product approval
Artificial intelligence and machine learning	What a performance claim is bound to before it can be read at all, and where external validation sits in that list	The step from performing to deployable stated as a step, and no performance figure of any kind
Biomarker-driven and computational psychiatry	Analytical validity, clinical validity and clinical utility separated, with the evidence each one actually needs	The rungs kept in order, and the reason a marker result licenses no conclusion about a patient
Digital therapeutics and prescription apps	What a prescription digital therapeutic is as a regulatory object, why it is a device and not a drug, and what engagement and attrition do to an effect that a tablet does not have	The blindness of this chapter's instruments stated on your own authority, no regulatory verb on any named product, and deployment left to the Community psychiatry and public health notes

21.3 Six prompts for the individual agents

These sections are compact and are marked on the same two things: what the receptor signature commits the agent to, and what each regulator's record says. Write each as cells and a pharmacology, not a paragraph about a drug.

- Asenapine: the receptor signature, the formulation constraint that follows from it, and then, cell by cell, which regulators hold a grounded record and what each one's wording covers.
- Lurasidone: the metabolic and movement profile and the indications, then the reason no first Indian approval date appears in this chapter, and what the Indian corpus does hold for it instead.
- Brexpiprazole: the receptor signature and the indications it carries, and what its United States record teaches about a citation whose fields do not all point at the same thing.
- Cariprazine: the receptor signature that separates it within its own class, the indications, and what a verbatim quotation loses when it is cut at a fixed character count instead of at a clause.
- Lumateperone: the receptor signature and the indications, then the jurisdiction whose record is supportive-source-only rather than grounded, and the sentence that record does not license.
- Zuranolone as an agent profile: the compact agent fields alone, the neurosteroid case left where it was argued, and what a profile may carry that a narrative section already owns.

21.4 A second pass, sorted by the way an answer fails

Once the eighteen prompts run cleanly, run the chapter in the other direction, sorted by how an answer goes wrong rather than by topic, because a supplementary question arrives as a defect and not as a heading.

HOW THE ANSWER FAILS	THE PROMPT THAT TRAINS IT	WHERE IT IS TESTED
A status carried across a border	Take any agent this chapter names and write its cells one at a time, refusing to let a full cell fill an empty one	The esketamine, orexin and endoxifen sections, and the agent prompts above
An answer given at the wrong grain	For every agent whose cells do not read alike, say why a single agent-level verdict cannot be written at all	The orexin and endoxifen sections
A negative built out of a silence	For each gap, say which instrument was pointed at it, how sensitively, and whose authority a claim would finally rest on	The muscarinic, psychedelic and digital-therapeutic sections
A recited approval read as a granted one	Name the state that is neither a grounding nor a gap, and the sentence it licenses in place of the obvious one	The lumateperone section
A figure supplied where no lane holds one	For each evidence-driven section, list the quantities that are unavailable and say why writing one would be a fabricated citation	The neurosteroid, rapid-acting, serotonin-hypothesis, psychedelic, machine-learning and biomarker sections

21.5 Where these marks are actually lost

Three silences in this chapter look identical on the page and license entirely different sentences: one where a named instrument was pointed at the question and found nothing, one where the instrument could not have answered whatever the truth was, which is **instrument blindness**, so its quiet carries no information, and one where another chapter holds the fact and this one declines to restate it so two accounts cannot drift apart, which is a **delegated gap** and the only one of the three that is not an absence of fact. Say which of the three a silence is before writing anything from it; the section owning the third case works the discrimination out in full.

The other loss is quantitative. Every evidence lane behind this chapter is unrun, so no efficacy figure, effect size, response rate, trial name, sample size, endpoint, onset time, half-life, dose or adverse-event rate is available to any section in it. Supplying one from memory is not a weak answer but a **fabricated citation**, which nothing said afterwards can repair. Rehearse the substitution: say what is known qualitatively, name the ordering, and name what would have to be read for a quantity to be stated.

TRAP **Mistaking recognition for recall.** A section that reads as familiar on a second pass has been recognised, not retrieved, and recognition collapses under the first supplementary question. A prompt that withholds its answer leaves nothing to recognise. If a prompt cannot be written out in full sentences, with a jurisdiction on every status clause and a named source beside it, it is not yet known, however familiar the section felt.

21.6 The Indian limb, rehearsed as a separate answer

Write the Indian part of any answer separately, and last, because it is the limb most often filled in from whichever record the candidate has read.

IN INDIA

Write the Indian limb from an Indian document or do not write it. The authority is CDSCO and the Drugs Controller General of India, and a committee minute, a clinical-trial no-objection certificate, a CTRI registration, a brand on a pharmacy shelf and a manufacturer's own copy are none of them a marketing authorisation. Where no Indian record was located, the answer that earns the marks names what was searched, on what instrument, and whose authority a claim would rest on, and stops there rather than hardening into a categorical negative, because the corpus behind these notes carries a measured limit and an absence in it is evidence and not proof. Then take the question the examiner is actually asking: what an Indian service must establish first, meaning who confirms the diagnosis, who owns the decision, who carries the follow-up, and what the patient can afford.

21.7 What to carry into the paper

Consolidation buys faster access rather than more material, so the last week on this chapter should look nothing like the first. Four habits, in order.

- Produce rather than review. Write one section out cold each sitting and mark it against its source, instead of rereading three and feeling three times as prepared.
- Write the cells before the pharmacology. If the status will not come, no amount of receptor detail rescues the answer.
- Rework each vignette with the sections shut, because a vignette is where a border crossing surfaces as a decision rather than as a sentence, which is the form the examiner will use.
- Close every answer by naming what you hand to another chapter, or, where this chapter owns the section in full, by saying so and giving the full account.

What survives when the detail goes is short. Most questions here are about a record, so name the regulator before the agent; where a section has no regulatory record, name the construct and its limit instead. Every silence is a statement about an instrument, so say which instrument and how sensitively it looked. Every quantity is unavailable, so give the ordering and name what would have to be read. Every boundary is a sentence handed over rather than material rebuilt. Those four moves let a candidate answer on a section they revised badly.

Consolidating this chapter means producing it and not rereading it: one section at a time, cold, written as cells before pharmacology, with every status clause carrying its regulator and its source, and every silence carrying what was searched.

PART X

X

Community Psychiatry and Public Health

Psychiatry at the scale of a population: the burden, the treatment gap, the programmes meant to close it, and prevention that is more than advice.

01 National programmes and mental health legislation

02 Epidemiology, burden and treatment gap

03 Suicide prevention and farmer's suicide

04 Stigma and anti-stigma interventions

05 Psychiatric rehabilitation and community care

06 Primary care, integration and service delivery models

07 Special-setting and disaster mental health

08 Prevention, economics and digital mental health

09 Culture, global mental health and traditional healing

10 Consolidate and apply

VISUAL 10.1

India's public mental-health stack



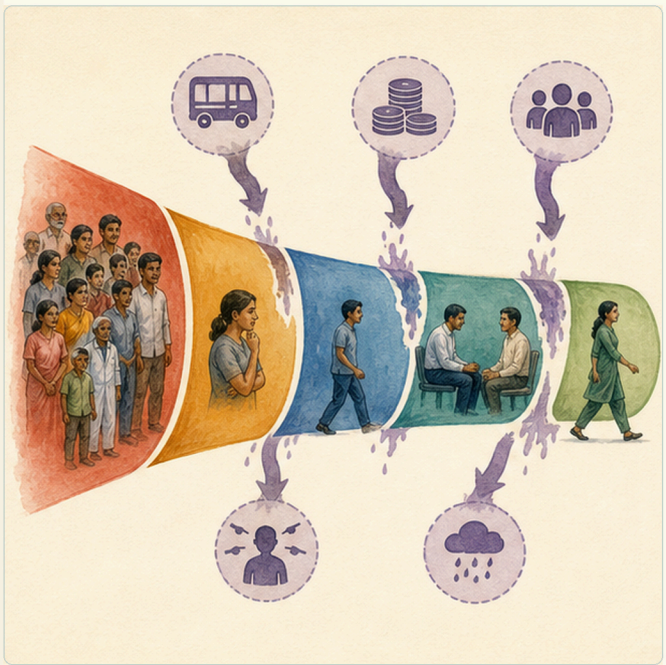
- National policy
- District services
- Primary care
- Tele-mental-health
- Rights and legal protection
- Community and family

Public mental-health care depends on linked policy, district implementation, primary care, tele-services, rights protection and community support.

Source lock: day-39.final.html#sec-t09

VISUAL 10.2

The treatment gap is a leaking pathway

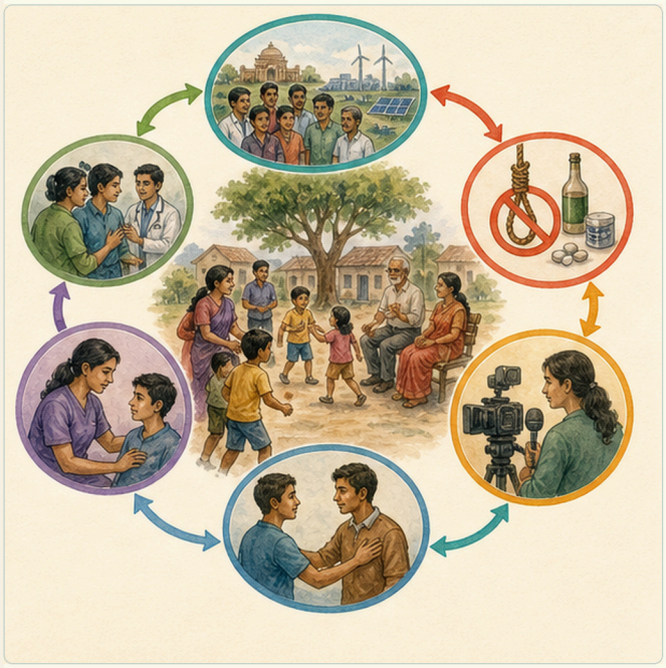


- Need
- Recognition
- Access
- Treatment
- Continuity
- Structural barriers

The treatment gap accumulates across recognition, access, initiation and continuity, shaped by structural barriers rather than a single point of failure.

Source lock: day-39.final.html#sec-t11

VISUAL 10.3



Suicide prevention works across multiple layers

- Population policy
- Means restriction
- Responsible media
- Gatekeepers
- Urgent clinical care
- Aftercare

Suicide prevention combines population measures, safer environments, responsible communication, gatekeepers, timely clinical care and aftercare.

Source lock: day-39.final.html#sec-t16

VISUAL 10.4



Rehabilitation connects treatment to citizenship

- Person-defined goals
- Family and community
- Skills
- Education and employment
- Clinical care
- Participation and recovery

Psychiatric rehabilitation links clinical care to skills, relationships, work, participation and person-defined recovery.

Source lock: day-39.final.html#sec-t29

PAPER IV

Community psychiatry and public health

Nine bands . one complete notes document

National programmes and mental health legislation

1 · Community psychiatry and public health: what this chapter owns, and what it points to

A programme that exists on paper is not a service, and a service that exists in a district is not yet care that a patient has received. The distance between those three things is the subject of this chapter. The national programmes, the money that moves them, the routes by which a person with an illness arrives at somebody able to treat it, and the evidence that any of it worked are all statements about a system rather than about a patient. That change of unit is easy to describe and hard to hold under examination pressure, because every other chapter in this paper has trained you to answer at the level of the person sitting in front of you.

So it is worth stating plainly what this chapter is for: programmes, financing, access pathways, service architecture, population interventions and their evaluation. It is not a second course in mental health law, not a course in professional ethics, not a course in bedside consultation, and not a course in epidemiological method, although it borrows vocabulary from all four and is examined alongside them. The marks here are given for saying what a health system does, at which level it does it, with what resources, for which population, and how anybody would know whether it made a difference. A candidate who answers that with the provisions of an Act, or with the assessment of one patient, has written a good answer to a question nobody asked.

1.1 The lens: the programme, the service and the population

The organising move is a change of unit. Clinical psychiatry takes the individual patient as the unit of analysis and asks what is wrong and what should be done about it. **Public-health implementation** takes the population and the service as its unit and asks a different set of questions: how many people in a defined group carry a condition, how many of them reach treatment, where along the route the remainder are lost, who is expected to act at each level of the health system, and what evidence exists that the arrangement changes anything at all. The clinical facts do not disappear when the lens changes. They stop being the answer and become inputs to a question about delivery.

Two constructs carry most of the weight and recur in every band ahead. **Service architecture** is the arrangement of who does what at which level: the community health worker, the medical officer in a primary care facility, the district team, the general hospital psychiatry unit, the state mental health institution and the specialist tertiary centre. **An access pathway** is the route an individual actually travels, from the first recognition that something is wrong to the point where an effective treatment is delivered, including the long stretches of that route that lie outside the

health system entirely. Almost every intervention taught in these notes acts on the architecture or on the pathway, and deciding which is usually the first useful thing to say about it.

RULE **The answer is written at the level of a population or a service, not at the level of one consultation.** Before answering any question in this territory, say which population is being described and which level of the system is expected to act. A clinically impeccable account of how a patient should be assessed and treated earns very little here, because the question was about who reaches that assessment, who is available to deliver it, and who pays for it. Keep the clinical content in a clause and spend the answer on the delivery decision.

1.2 What this chapter deliberately does not teach

These notes border four others closely enough that material leaks in both directions, and each border is worth rehearsing as a sentence. Naming the owning chapter accurately is faster and better rewarded than rebuilding its material from memory, and it is the only way an examiner can see that you knew where the boundary was.

- The statutes themselves are not taught here. Their provisions, the machinery they create, the rights they confer, the safeguards they impose and the procedures they prescribe belong to the Mental health legislation and forensic psychiatry notes. What belongs here is the service consequence: what a law obliges a system to provide, who has to be in post before that provision can exist, and how it is financed and counted.
- Professional conduct, capacity and consent, custodial procedure and the statutory frame governing disaster response belong to the Forensic ethics, consent and legal procedure notes. These notes take the same events from the service side: how a response is organised, who staffs it, at which phase, and what happens to routine care while it runs.
- The bedside act belongs elsewhere as well. The cultural formulation of an individual case, the idioms in which a particular patient describes distress, the named culture-bound presentations and the clinical assessment of a patient already under the care of a traditional or faith healer are taught in the Consultation-liaison, geriatric and transcultural psychiatry notes. This chapter takes the same ground at population level: help-seeking, the sequence of providers a community actually consults, and the relationship between the health system and the other places people go.
- Method belongs to the Research Methods, Biostatistics and Epidemiology notes: what incidence and prevalence are as measures, how a survey is designed and sampled, what a study design can and cannot support, and how a rate is standardised. These notes consume the products of that method rather than teaching it, and the discipline they add is about reading a published figure correctly rather than about producing one.

Two narrower handovers are worth the same treatment. The clinical assessment and management of the individual patient at risk of suicide belongs to the Somatic-symptom, sexual, impulse-control disorders and suicide notes, while the population-level prevention of suicide is taught here in full. The clinical care and protective assessment of the individual child belongs to the

Child psychiatry: assessment, treatment, services and protection notes, while the service and programme interface with child protection law sits here.

TRAP **Answering a programme question by reciting the statute.** Asked how mental health services are organised in a district, candidates produce the sections and safeguards of the governing legislation, because that material is adjacent enough to feel responsive and was revised harder. The recitation may be accurate and still earn little, because the examiner asked what is delivered, by whom and at what level. State the legal obligation in a clause, then spend the answer on the structure built to discharge it.

1.3 What makes a national figure true

This is the most number-heavy territory in the paper, and almost every figure in it is one you have heard quoted at a seminar. That is exactly the problem. A national figure carries meaning only when four things travel with it, and a figure quoted without them is not a fact but a rumour with a decimal point.

Source	Reference year	Population	Denominator
WHICH BODY ISSUED IT	THE PERIOD IT DESCRIBES	WHOM IT COUNTS	WHAT IT IS A PROPORTION OF

Each field fails in its own way. Drop the source and the number rests on the speaker, which under examination means on memory. Drop **the reference year** and a survey that describes one period is offered as a description of the present, which matters more here than anywhere else in the paper because the national instruments in this field are re-run rarely and their findings stay in circulation long after the period they measured. Drop the population and the commonest error in the whole territory follows: the same survey's finding for one group is quoted as though it applied to another, so a figure that is perfectly accurate becomes a false statement about the wrong people. Drop **the denominator** and a count is read as a rate.

The last field needs one remark, because the honest answer is sometimes that no valid denominator exists. A total of sanctioned units, a cumulative count of calls to a helpline and a national count of deaths in an occupational group are all real numbers with no meaningful proportion attached, either because nobody counted the base or because the base is not knowable. The correct sentence says so, and stops: a cumulative count of calls is a count of calls, never a count of people. Refusing to supply a denominator is not weakness in an answer. It is the only version of the sentence that is true.

PITFALL

Reading a sanctioned figure as a delivered one. Programme documents in this field count what has been approved, funded or covered on paper, and those counts are published, quotable and entirely genuine. They are not counts of functioning services. A district may appear in an approved list with posts unfilled, and a facility may be listed with no specialist attached. The same trap runs through activity data: a service that reports contacts is reporting contacts, not distinct individuals, and certainly not individuals who completed treatment. Before quoting any coverage figure, ask what the noun beside the number actually is, and say that noun out loud in your answer.

1.4 The map of the territory

Nine bands follow. They are set out below not as a list to memorise but as a routing table: given a question, it says which band holds the answer. The bands themselves supply the programmes, the figures, the models and the attributions.

BAND	WHAT IT ESTABLISHES	THE KIND OF ANSWER IT REWARDS
National programmes and the legislative frame around them	What the state has undertaken to provide, through which programme structures, and at which level of the system each acts	Components and objectives, the level each one operates at, and the difference between what is sanctioned and what is running
Burden and the treatment gap	How much illness a defined population carries, and what proportion of it reaches treatment	A figure with its source, its reference period, its population and its denominator attached, and the determinants behind the shortfall
Suicide prevention at population level	Which population-level actions bear on deaths by suicide, and which groups carry raised risk	Actions classified by the population they address rather than by the individual they treat, with the level of action named
Stigma and the interventions aimed at it	How stigma operates in the public, in the person and in institutions, and what has been tried against each	The type of stigma identified first, then the intervention matched to that type and the evidence behind it
Rehabilitation and community care	What happens after the acute episode, and where a person then lives, works and is supported	The components of a rehabilitation package and the setting each belongs to, with residential and daytime options distinguished
Primary care, integration and delivery models	Who delivers mental health care where no specialist exists, and what supports them in doing it	The task each cadre performs, the supervision structure behind it, and the model that organises both
Special settings and disaster	How services are organised for populations that ordinary outpatient care does not reach	The constraints the setting imposes first, then the service response adapted to those constraints
Prevention, economics and digital care	What can be prevented, what illness costs, how care is paid for, and what technology changes about reach	The level of prevention named, the economic question stated precisely, and a claim about reach kept separate from a claim about benefit
Culture, global mental health and traditional healing	How explanatory models and existing sources of help shape the route into care	The pathway described as it is actually travelled, and the working relationship between the health system and other providers of help

1.5 The spine underneath the map

The bands are not nine separate topics; they are one argument told in sequence. Burden establishes the size of the problem in a defined population. The treatment gap converts that burden into a statement about the service, because a gap is not a fact about illness at all but a measure of how badly access has failed. The pathway work then explains where along the route people are lost, and the architecture says who might have caught them at each of those points. The intervention bands, from suicide prevention through stigma, rehabilitation and primary care integration, are all attempts on one or the other: they change who can act, or they change how a person reaches somebody who can. Prevention and economics ask which of those attempts is worth making and who pays for it. Evaluation asks whether any of it moved.

Read that way, the last band is not an appendix. Culture, explanatory models and existing healers sit at the beginning of the pathway rather than at the end of the chapter, because for a large part of the population the first consultation for a mental health problem is not with the health system at all. A reader who meets that band last but places it first has understood the sequence.

1.6 Inputs, outputs and outcomes

Evaluation is the part of this territory candidates most often skip, and it is where an understood answer separates fastest from a memorised one. Programme evidence comes in three kinds, and confusing them is the single most productive source of wrong statements in the whole field.

WHAT THE FIGURE COUNTS	A QUESTION IT CAN ANSWER	A QUESTION IT CANNOT ANSWER
An input: what was committed, sanctioned, staffed or funded	Whether the system undertook to act, and on what scale	Whether anything was delivered to anybody
An output: what the service did, counted as activity	Whether activity took place, and how much of it	Whether the people who most needed the service were the ones who got it
An outcome: what changed in the population	Whether the health of a defined group moved, and in which direction	Which part of the arrangement produced the change, unless the evaluation was designed to attribute it

The distinction has a companion in the idea of **functional coverage**, which asks not how many units exist but how many are staffed, supplied, open and used by the population they were built for. Sanctioned coverage is an input; functional coverage sits between an input and an output; and neither is an outcome. Where the functional figure has not been measured, the correct statement gives the approved figure, says that it is an approval, and declares the working figure unavailable. That declaration is a complete answer, and a better one than a confident number that quietly promotes a sanction into a service.

GOOD TO REMEMBER

Before quoting any national figure in a case presentation, a ward round, a teaching session or an answer script, ask four questions in order. Who issued this number. What period does it describe, as distinct from when the document was published. Which population does it count, and is that the population I am about to apply it to. And what is it a proportion of, or is it a bare count with no valid denominator. A clinician who runs those four before speaking will not be the person who quotes a survey's finding for one group as though it described another, which is the most frequently repeated error in this whole field.

1.7 What a good answer in this territory looks like

A small number of moves earn most of the marks here, and they recur across all nine bands in different clothing. They are worth learning as habits, because the detail here changes while the moves do not.

- Naming the level of the health system at which the answer operates, before describing what happens there. Community, primary care, district, general hospital, state institution and tertiary centre are not interchangeable settings, and most questions here are really asking which one you would use.
- Separating what a policy or programme undertakes from what a service delivers, and saying which of the two the evidence you are quoting actually measured.
- Attaching the source, the period, the population and the denominator to every figure, and declining to give a figure at all where any of the four is unavailable to you.
- Handing over what a neighbouring chapter owns in a single sentence that names it, rather than reconstructing statute, ethics, bedside method or epidemiological technique from memory in the middle of a service answer.

Those four are also, in the same order, the four ways an answer here goes wrong. An answer written without a level reads as an opinion about mental health in general. A figure without its scope is a wrong number wearing a citation. And an answer that drifts into the law has left the marks on the table while writing accurately about something else.

1.8 A scaffold to carry into essays and vivas

Under examination pressure the detail goes first, and in this field the detail is exactly what is hardest to hold, because it consists of programme names, structures, reference periods and figures that resemble one another. What has to survive is a routing sequence. The scaffold below is a teaching device rather than a validated framework, and it will structure an answer on any band in these notes, from a programme to a delivery model to a question about culture and help-seeking.

MNEMONIC

Population, Pathway, Programme, Proof

Population: name the group the question is about and, if a number is coming, what it is a proportion of. **Pathway**: describe how that group actually reaches care and at which step it stops, including the steps taken outside the health system. **Programme**: name the level of the system expected to act, what has been put there to act with, and who supervises it. **Proof**: say what would count as evidence that it worked, and whether the figure you have is an input, an output or an outcome.

For a short answer, run the four and stop. For a long essay, expand each into a paragraph and take the specifics from the band that owns them. In a viva, make the sequence audible: say whom you are talking about, how they reach care, who is meant to be there, and how you would know whether it helped. That last sentence separates a candidate who has learned a list of programmes from one who has understood why they exist, and it is also what a district team needs from the psychiatrist it consults.

A programme is not a service, a service is not access, and access is not care: name which of the three a question is asking about, and attach the source, the period, the population and the denominator to every figure you offer.

2 · National Mental Health Programme (NMHP): components, objectives and evolution

A national programme is not a policy, not a statute and not a scheme, and the marks in this area are lost by candidates who cannot tell the four apart. The **National Mental Health Programme** is the instrument through which the Union government made mental health a delivery obligation of the general health system rather than the business of a handful of specialist hospitals. Everything else in Indian community psychiatry hangs off it: the district limb that carries services to the population, the training and workforce apparatus, the money that reaches a state, and the reporting through which any of it becomes visible. Understand what kind of object it is and the rest of this chapter becomes a set of consequences.

This section teaches the programme as programme architecture: what it is for, what it consists of, how it has moved, how it is financed and monitored, and what can honestly be said about how far it has reached. Its district implementation is a separate section of this chapter. Its role as the delivery frame for community rehabilitation of a person with a psychotic illness, set against assertive community treatment, is taught in the Schizophrenia: management, antipsychotics and adverse effects notes, and is not restated here. The statutes the programme now runs inside are taught in the Mental health legislation and forensic psychiatry notes.

2.1 Four instruments, four different kinds of authority

Before the programme, the vocabulary. Indian mental health administration uses four words that examiners deliberately mix, and each carries a different kind of force.

INSTRUMENT	WHAT IT FIXES	WHO IS BOUND	HOW FAILURE SHOWS
Statute	Rights, duties, offices, procedures and remedies	Everyone within the jurisdiction, on pain of legal consequence	An entitlement exists and is not honoured; the remedy is legal
Policy	Direction, priorities and the populations named as vulnerable	Government, as a statement of intent	The direction is stated and nothing is funded to move in it
Programme	Objectives, components, delivery architecture and reporting	The health system that operates it	Approved on paper, unstaffed in practice
Scheme	A defined benefit or activity with its own budget line	The implementing agency and the beneficiary	Money released, benefit not received

The programme sits in the third row, and its characteristic failure mode is the third one. That single row is the most examinable idea in this section and it recurs in every part of this chapter: an approved unit is not a functioning unit, a sanctioned post is not a filled post, and a released instalment is not a service delivered. Hold the distinction and you will not misread a coverage figure again.

RULE **A programme is an architecture plus a reporting system, and it can only be judged on what it reports.** Its objectives tell you what it was designed to change; its components tell you the machinery it intends to use; its monitoring tells you what evidence of change will ever exist. Where the third is thin, no amount of the first two will let you say whether the programme worked.

2.2 Objectives, and the one this chapter can put a source against

The objectives of this programme are among the most recited facts in Indian psychiatry and among the least reliably sourced. The founding programme document, the one that sets out the original objectives, could not be retrieved from a serving official source during the preparation of this chapter: the ministry's own page returned nothing after a domain migration, a repair attempt returned the same dead address, and no archived snapshot was available. That is stated here rather than papered over, because the honest position is itself an answer. The objectives as they are usually recited in teaching may well be accurate; this chapter simply cannot certify them against a document it could open, and it will not print a list it cannot stand behind.

What can be certified is the objective the programme's current official page carries. The Directorate General of Health Services states an objective of the National Mental Health Programme in these terms: **to enhance human resource in mental health sub-specialties**. That is a live official statement, read from the programme's own live page, and it describes the mental health subspecialty workforce of India. It is not a rate and it has no denominator, which is correct and not a defect: an objective is a direction of travel, not a proportion of anything.

Notice that this objective is itself evidence of evolution, which is the second half of what this item covers. An objective about human resource appearing on the current page tells you that the

programme, whatever it was designed to do at the outset, now regards the production of specialist and subspecialist manpower as one of its own deliverables rather than as something the education system supplies to it from outside.

IN INDIA

Because the programme names enhancement of human resource in mental health sub-specialties as an objective of its own, a department planning against it should treat training capacity as a programme output to be planned and reported, not as an academic by-product of clinical work. That changes what a unit counts: not only patients seen, but trainees trained, supervisors available, and the subspecialty posts that a trained person can actually move into. A workforce objective without a post at the end of it produces qualified people for other systems.

2.3 Why human resource is the binding constraint on this programme

Public health programmes usually fail on one of four constraints: money, materials, people or information. Mental health services are unusual in that the first two are comparatively cheap. The essential medicines are old, off patent and inexpensive; the physical infrastructure is a room, a record and a reliable supply. What the service cannot substitute for is trained people, because the intervention *is* a trained person's time, whether that person is a psychiatrist, a clinical psychologist, a psychiatric social worker, a nurse or a trained medical officer at the first level of care.

That has three design consequences which run through the rest of this chapter.

- **Capacity is built slowly and lost quickly.** Specialist training takes years, while a transfer, a resignation or an unfilled sanctioned post removes the capacity immediately. A programme that expands its coverage faster than it produces people converts a staffing problem into a coverage illusion.
- **Task-shifting is not optional, it is arithmetic.** If the specialist workforce cannot be enlarged at the rate the population needs, some part of the work must move to a broader cadre with supervision. The mechanism, the guidance behind it and the medical officer's role in it are taught in their own sections of this chapter.
- **Retention is a programme variable.** Where a trained person is posted, whether the post is at the level they were trained for, and whether supervision exists, determine whether training converts into service. None of that is visible in a count of trainees.

2.4 The architecture, from the national programme down to the district

A national programme in India is not delivered nationally; it is delivered by states, in districts. Health is administered at state level, so the Union instrument sets objectives, defines components, provides a share of the money and asks for returns, while the actual service is run by a state health department through a district unit. That two-level structure explains most of what is puzzling about Indian programme data. The centre knows what it approved and what it released; the district knows what it staffed and what it delivered; and the two are not the same fact.

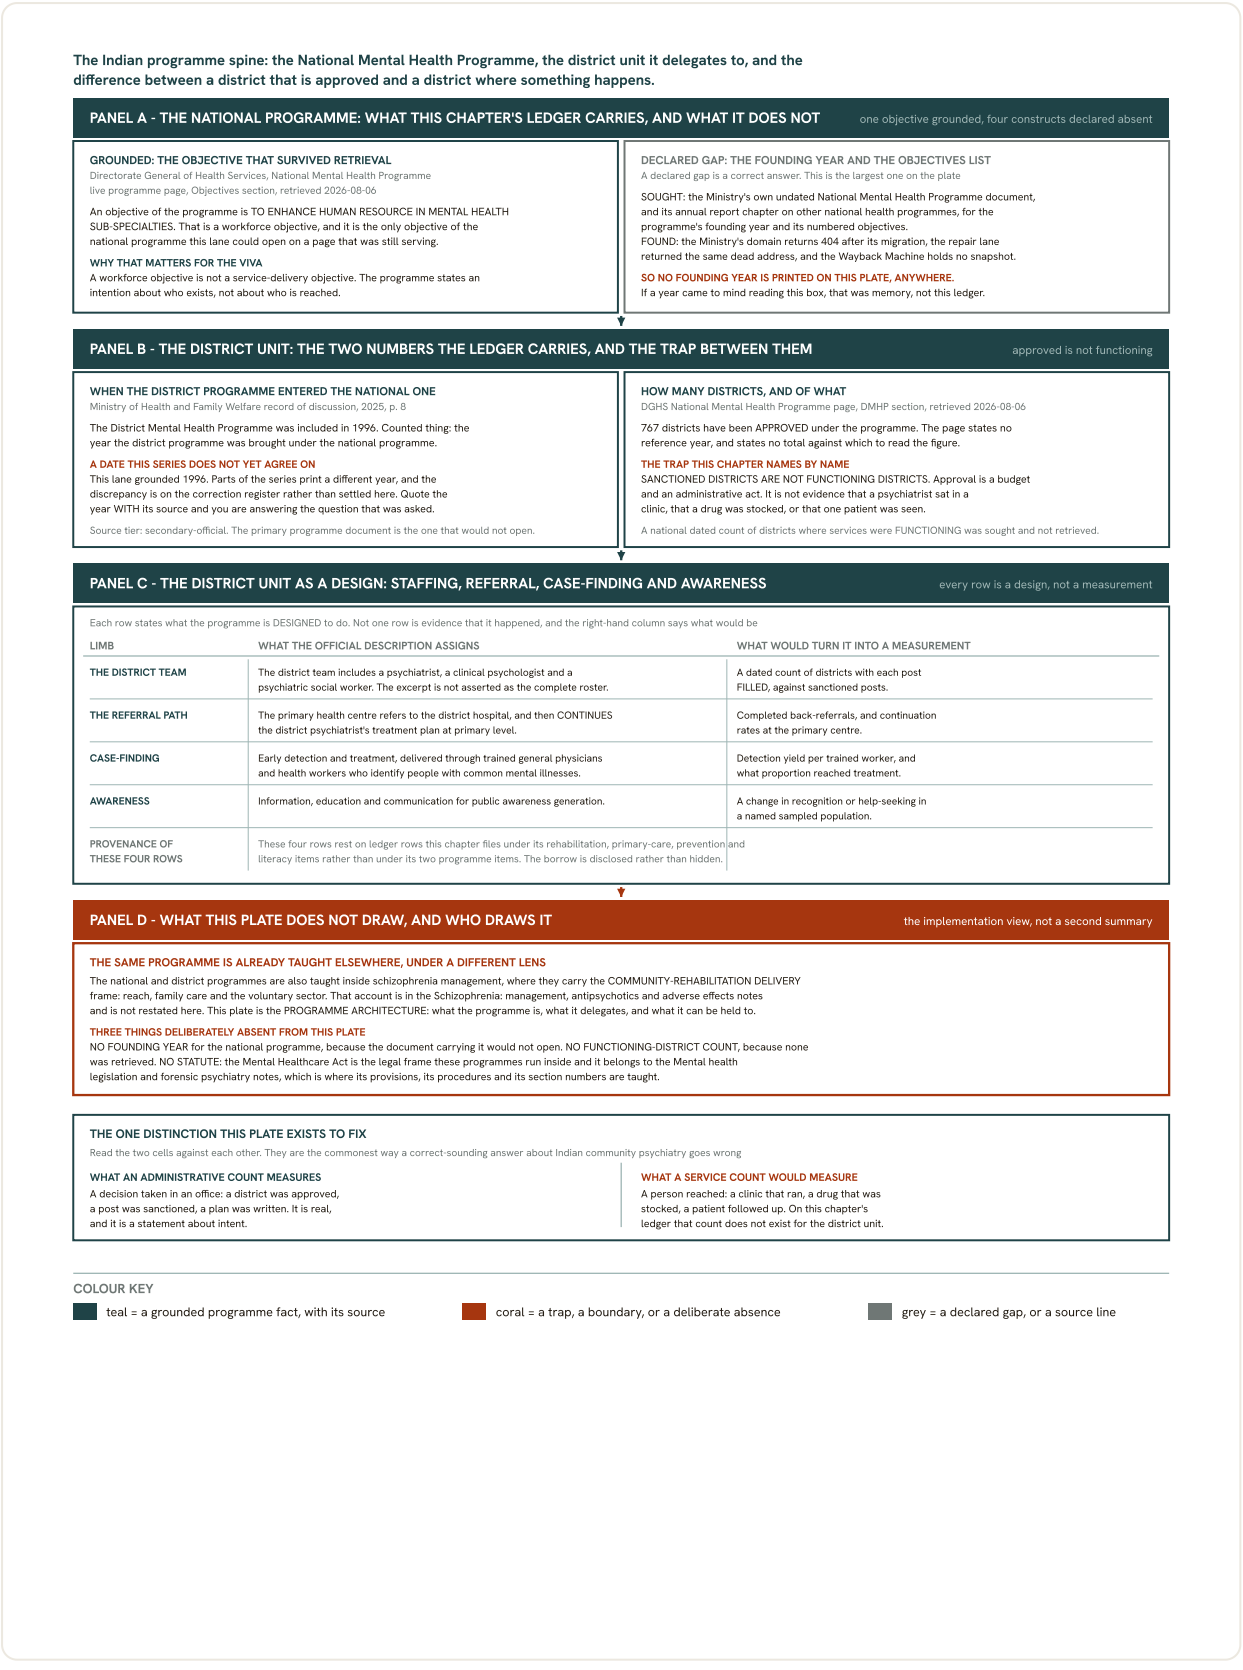


Fig 39.1 The Indian programme spine, drawn as an implementation view. The national programme contributes one retrievable objective, to enhance human resource in mental health sub-specialties; its founding year and numbered objectives are a declared gap, because the Ministry's own documents no longer serve and no archived snapshot exists, so no year is printed. The district programme was included in 1996 and is reported as approved in 767 districts, with no reference year and no stated total. The chapter's own trap sits between those two facts: approval is an administrative act, and a national count of districts where services were functioning was sought and not retrieved. The staffing, referral, case-finding and awareness rows state design, never delivery.

The accompanying figure sets out that spine, from the national programme down through the district unit to the components, staffing and referral arrangements at the point of delivery, and marks the boundary where the community rehabilitation account in the Schizophrenia: management, antipsychotics and adverse effects notes takes over. The district limb, the District Mental Health Programme, has its own section immediately after this one and its structure, components, staffing and coverage evidence are taught there rather than here.

What belongs here is the logic of the relationship. A national programme can do four things and only four: it can set an objective, define a component, move money, and require a return. Everything else, including whether a post is filled, whether a drug is on the shelf and whether a patient is seen, happens at a level the national instrument can influence only through those four levers. When you are asked why a national programme underperforms, the answer is almost always that one of the four levers is weak, and most often it is the fourth.

2.5 Evolution, and the document that could not be opened

The programme has moved considerably since it began: from an initial concentration on integrating mental health into general health care, through the addition of a district implementation limb, towards a wider portfolio that now includes workforce production, tertiary capacity and digital service delivery. The programme's own annual reporting is where those movements are recorded plan period by plan period, and it is the document a candidate should ideally cite.

Two of this chapter's declared evidence gaps sit exactly there. The ministry annual report chapter that records the programme's position among the other national health programmes could not be retrieved from a serving official address: the domain returned nothing after migration, a repair pass returned the same dead address, and no archived copy was available. Two separate records in this chapter's evidence ledger point at that same unreachable section. The consequence is that this chapter describes the direction of the programme's evolution, which is visible from its current official page and from the instruments that surround it, and prints no dated milestone list, because the document that would date the milestones could not be opened.

TRAP **Writing the year a programme started from memory.** Programme years, policy years and Act years are the most quoted and most misquoted facts in this whole area, and they feel exactly as certain from the inside as a year you have actually read in a document. This chapter prints a programme year only where a source it could open carries one. Where it could not, it names the programme without the year and says so. In a viva, "the founding document is not currently retrievable from an official source, so I would not quote a year I cannot cite" is a stronger answer than a confident wrong date.

2.6 How you would actually judge this programme, and what a figure must carry

Judging a national programme means choosing a denominator, and most published commentary skips that step. There are four honest questions and they get progressively harder to answer.

- **Reach.** How many districts, units or facilities are covered? This is the easiest to answer and the easiest to misread, because approval and function are recorded by different people for different purposes.
- **Function.** Of the covered units, how many are staffed, supplied and open? This is the number that matters and the number most often missing.
- **Contact.** Of the people in the population who need care, how many make contact with the service? This requires a population estimate as a denominator, which comes from a survey rather than from the programme.
- **Outcome.** Of those who make contact, how many are adequately treated and how many improve? This requires a follow-up system, which is the weakest layer of Indian programme reporting.

Source	Year	Population	Denominator
WHICH BODY ISSUED THE FIGURE	WHICH YEAR IT REFERS TO, NOT WHEN PUBLISHED	WHICH PEOPLE IT DESCRIBES	WHAT IT IS A PROPORTION OF

Those four tiles are the discipline this entire chapter runs on, and they are why so few programme figures appear in it. A national number missing any one of them is not a fact; it is a recollection with a decimal point attached. Where a programme figure carries all four, this chapter prints it with all four. Where it does not, the chapter names the construct without its number and records the absence, which is the correct answer and not an evasion.

MNEMONIC

Objectives, components, delivery, money, returns

The five things any national health programme consists of, in the order you should describe one. An answer that gives the objectives and stops has described a fifth of the instrument.

PITFALL

Describing the programme entirely through its objectives. Objectives are the most memorable part and the least informative, because every national programme's objectives read approximately alike. What differentiates one programme from another is its delivery architecture, the level at which the money stops, and what it requires anyone to report. If you can describe those three for the National Mental Health Programme, you can answer almost any question in this area, including ones about programmes you have never studied.

The National Mental Health Programme is a Union instrument that sets objectives, defines components, moves a share of the money and requires returns; states and districts do the delivering, which is why approval and function are never the same number.

3 · District Mental Health Programme (DMHP): structure, components and implementation

The district is where Indian mental health policy either becomes a service or stays a sentence. The **District Mental Health Programme** is the operational limb through which the National Mental Health Programme reaches a population: a team, a budget line, a place to work from, a referral route in both directions, and a training obligation towards the general health workforce of the district. Every national ambition in this chapter, integration into primary care, task-shifting, closing the treatment gap, reaching a rural population, is discharged at this level or nowhere.

This section teaches the district limb as a service architecture and as an implementation problem. The parent programme's objectives, evolution and national machinery are taught in the section immediately before this one and are not repeated here. The way this same district structure functions as the community delivery vehicle for rehabilitation in schizophrenia, set beside assertive community treatment, is taught in the Schizophrenia: management, antipsychotics and adverse effects notes.

3.1 Why the district, and not the state or the block

The choice of the district as the unit is not administrative convenience; it is the level at which the three things a mental health service needs coincide. The district is small enough that a single specialist team can physically reach its facilities and large enough to justify one existing at all. It is the level at which the health administration already holds a budget, a stores system, a supervisory hierarchy and a reporting line, so a new programme can be attached to machinery that exists rather than built beside it. And it is the level at which secondary care sits, which matters because a mental health service that has no bed and no acute pathway becomes an outpatient clinic that refers its sickest patients away.

Compare the alternatives and the logic becomes obvious. A state-level unit can plan but cannot deliver; the distances defeat it and the specialists sit in the capital. A block-level unit can reach people but cannot be staffed with a specialist team at that density, and it has no bed. The district is the compromise, and like every compromise it inherits both problems in miniature: it is still far from the periphery, and it is still thin at the top.

RULE **A district programme is defined by what it does to the district's existing services, not by what it does itself.** Its clinic is the smallest part of it. Its real product is a general health system that can identify, start, continue and refer mental health care, which is why training and supervision, not consultations, are the components that determine whether it works.

3.2 When the district limb came in, and what that means

The **District Mental Health Programme** was brought in under the national programme in **1996**, and that year is stated in an official Ministry of Health and Family Welfare consultation record rather than in the founding instrument itself. That distinction matters and is worth carrying: the source is official but secondary, meaning it restates the fact rather than being the document that created it. A different year for the same event circulates widely in teaching

material, and this chapter prints the year its own retrievable official source carries rather than the year that is most often repeated.

The substantive point behind the date is that the district limb was an addition. The national programme began as an integration argument directed at the general health system; the district structure came afterwards as the vehicle for making that argument operational. That sequence explains a persistent feature of Indian mental health administration, which is that the district programme is always described as a component of something larger and is always funded as one, so its fortunes rise and fall with decisions taken about the parent programme rather than about itself.

3.3 What a functioning district unit actually consists of

Describe a district unit by its components and you will be able to say, for any real district, exactly which part has failed. Six things have to be simultaneously true.

- **A team.** A specialist core with psychiatry, clinical psychology, psychiatric social work and psychiatric nursing, plus programme and record staff. The composition matters less than the completeness: a team missing its psychosocial cadre becomes a prescribing clinic, and a team missing its record staff becomes invisible to the reporting system.
- **A place, at more than one level.** A base at the district hospital and outreach to peripheral facilities, because a service that exists only at the district headquarters has simply relocated the travel burden from the state capital to the district town.
- **An uninterrupted drug supply.** The essential psychotropics are cheap and off patent, so a stock-out is a logistics failure and never a cost one. A single interrupted supply undoes months of adherence work and teaches a family that the service cannot be relied on.
- **Referral in both directions.** Upward referral is easy to arrange and easy to overuse. Back-referral, the deliberate return of a stabilised patient to a peripheral facility with a plan the medical officer can follow, is the harder half and the one that determines whether the specialist team ever has capacity for new patients.
- **A training cascade.** The district team trains medical officers, health workers, counsellors and teachers, who then identify and manage what they can. This is the multiplier; without it the programme's reach is capped at the number of patients four or five specialists can personally see.
- **Records and returns.** Whatever is not recorded did not happen as far as the programme is concerned, and whatever is recorded without a denominator cannot be interpreted by anyone above the district.

Read that list again as a diagnostic instrument. When a district programme is described as non-functional, one of those six has failed, and the failure is usually the team or the supply. When it is described as functional but ineffective, the failure is almost always the fifth or the sixth.

3.4 The training cascade, and the school component that belongs elsewhere

The cascade is the part of the design that carries the public health logic. A specialist team that spends its time in direct consultation converts a scarce resource into a small number of individual encounters. The same team spending part of its time training, supervising and supporting a much larger general workforce converts the scarce resource into system capacity. That is the argument for task-shifting expressed at district scale, and the guidance and the medical officer's role in it are taught in their own sections of this chapter.

The programme's school outreach component is part of the same cascade, and it is taught in the Child psychiatry: assessment, treatment, services and protection notes rather than here. That chapter carries the school component in detail, including the teacher-training model with its own coverage figures for master trainers and schools. This section names the component as one limb of the district cascade and reproduces none of its numbers, so that a reader who meets a teacher-training figure meets it once, in one place, with one set of denominators.

3.5 Coverage: the number that exists, and exactly what it counts

There is one national coverage figure for this programme that can be printed with its provenance intact, and reading it correctly is the single most examinable skill in this section.

The Directorate General of Health Services reports that **767 districts** have been approved under the programme. That is approved coverage. It is not functioning coverage, and the source itself does not state the total number of Indian districts against which it should be read, so even the proportion approved cannot be computed from it without importing a denominator from somewhere else.

767	1996	Not stated
DISTRICTS APPROVED UNDER THE PROGRAMME, NOT DISTRICTS FUNCTIONING	YEAR THE DISTRICT LIMB CAME IN UNDER THE NATIONAL PROGRAMME	DISTRICT TOTAL, SO NO APPROVAL PROPORTION IS COMPUTABLE

Set the four fields against that figure and its limits are plain. The source is the programme's own official page. The population is districts of India considered for approval. The counted thing is districts approved. The denominator is not supplied by the source at all. Three of the four are strong; the fourth is missing, and the honest sentence says so rather than filling the hole with a district total remembered from elsewhere.

3.6 The number that does not exist, and why that is the answer

The obvious next question is how many of those districts have a service actually running in them. That figure was sought for this chapter from the national programme, from the directorate and from the ministry, including by targeted searches of official domains for functioning or operational districts, and by attempted retrieval of the ministry's own annual report. It was not obtained. There is, in this chapter's evidence position, no national, reference-dated, official count of districts in which district mental health services were functioning.

That absence is content, not a hole in the teaching. A candidate who says that the approved figure is the only national coverage number available, that approval is a sanctioning decision rather than

a service fact, and that the functional count is not published in a citable form has described the Indian evidence position accurately. A candidate who converts the approved figure into a claim about services running has made an error that no amount of confidence repairs.

COVERAGE QUESTION	WHAT WOULD ANSWER IT	WHAT EXISTS NATIONALLY
Is the district sanctioned?	An approval count from the programme's own records	Published: 767 districts approved under the programme
Is the service running?	A reference-dated count of districts with services functioning	Sought from the programme, the directorate and the ministry; not obtained
Is the team complete?	Filled posts against sanctioned posts, by cadre and by district	Not held in this chapter's evidence
Are people being reached?	Contacts against an estimated population needing care	Requires a survey denominator, taught in the survey section of this chapter

TRAP **Reading a sanctioned district as a functioning district.** This is the classic error in Indian community psychiatry and it is built into the way the data are produced: approval is recorded centrally because the centre approves, while function is a state and district fact that nobody aggregates. The approved count is real, official and correctly reported; it simply answers a different question from the one the reader is usually asking. Say "approved" every time you say the number.

IN INDIA

When a district is described to you as covered by the programme, treat that as an administrative statement and ask three follow-up questions before you plan anything on it: which posts in the team are filled today, whether the essential psychotropics were in stock last month, and whether any patient has been back-referred to a peripheral facility with a written plan. The published national figure is 767 districts approved under the programme, and approval is where those three questions begin rather than where they end.

3.7 What monitoring would have to look like

The reason the functional count does not exist is not indifference; it is that nobody has defined the unit of observation. Deciding whether a district service is functioning requires an operational definition that a clerk can apply consistently: how many days open, which cadres present, what stock position, what minimum activity, over what reference period. Without that definition, every state answers a slightly different question and the national total means nothing.

- An operational definition of a functioning unit, applied uniformly, with a stated reference period.
- Filled against sanctioned posts by cadre, because a team is not staffed if only the psychiatrist post is filled.
- Stock-out days for the essential psychotropics, which is the fastest single indicator of whether a service is real.

- Back-referrals as a proportion of stabilised patients, which measures whether the cascade is working rather than whether the clinic is busy.

MNEMONIC

Team, place, drugs, referral, training, returns

The six components of a district unit, and the six ways one fails. Name the six and you can diagnose any district service you are asked about, including one you have never visited.

PITFALL

Judging a district programme by its outpatient numbers. A busy district clinic can coexist with a completely failed cascade, and often does, because direct consultation is the easiest activity to increase and the easiest to count. Rising attendance at the district base with no back-referral and no trained medical officers means the specialist team has absorbed the district's workload instead of distributing it, and the service will saturate.

The published national figure is 767 districts approved under the programme; approval is a sanctioning decision, no national count of functioning districts is published, and the two must never be quoted as one.

4 · Mental Healthcare Act 2017: key provisions, advance directives, nominated representative

Every programme in this chapter runs inside a statute, and the statute is not this chapter's to teach. The Mental Healthcare Act is the legal frame within which India's national and district mental health programmes now operate. It fixes the entitlements a service has to honour, it creates instruments a patient may use to direct their own future care, and it establishes bodies that supervise both. Its provisions, including the advance directive and the office of the nominated representative, are taught in full in the Mental health legislation and forensic psychiatry notes. This section names the frame, prints that pointer, and stops.

That is a deliberate structural decision and not an omission. A statutory entitlement is a text with a fixed wording, and the fastest way to corrupt it is to let two chapters paraphrase it independently. The second account will differ from the first in a modal verb, a threshold or a procedural step, and neither reader will be able to tell which drifted. One owner, one wording.

4.1 The boundary, drawn once, and applied everywhere in this chapter

The division of labour is easier to hold if it is stated as a question rather than as a list of topics. A statute answers **what a service owes**: the right, the duty, the instrument, the procedure and the remedy. A programme answers **what a service has**: the posts, the drugs, the budget line, the referral route and the coverage. The two questions have different owners because they have different failure modes. A wrong entitlement misleads a patient about their rights; a wrong coverage figure misleads a planner about a population.

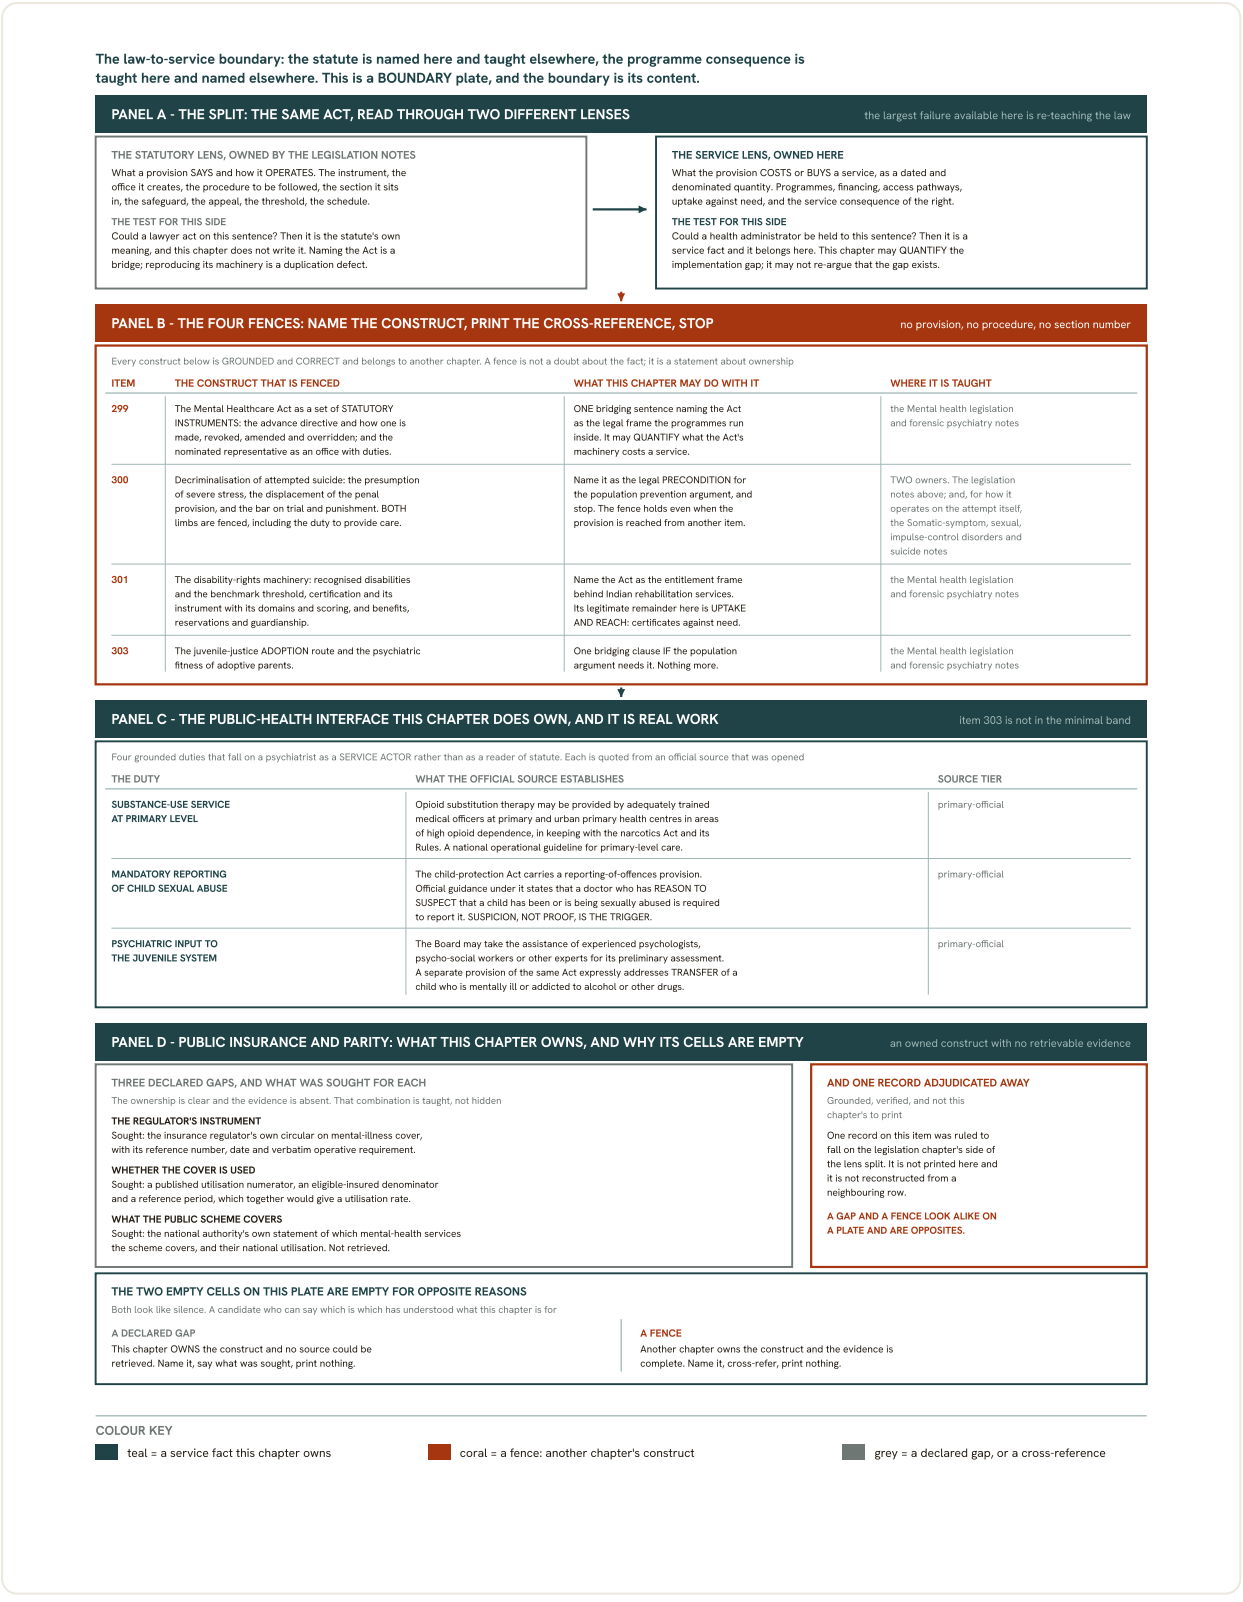


Fig 39.2 **The law-to-service boundary, drawn as an exclusion map.** The statute is named here and taught in the Mental health legislation and forensic psychiatry notes; what a provision costs or buys a service is taught here. Four constructs are fenced: the Act's statutory instruments, the decriminalisation of attempted suicide, the disability-rights machinery, and the juvenile-justice adoption route. No provision, procedure or section number is reproduced on this plate. The interface this chapter does own is real work and it is grounded: opioid substitution therapy provided by trained medical officers at primary level, a doctor's duty to report suspected child sexual abuse on reasonable suspicion rather than proof, expert psychological and psycho-social assistance to the juvenile-justice Board, and a transfer route for a child who is mentally ill or addicted. The insurance and public-scheme cells are empty because the regulator's own instrument could not be retrieved, which is the opposite kind of emptiness from a fence.

The accompanying figure is the exclusion map this chapter applies whenever a statute appears in it. Everything on the legal side is named here and taught elsewhere. Everything on the service side, meaning programmes, financing, access pathways, service architecture and population consequences, is taught here in full. No arrow crosses the line without a pointer attached to it.

RULE **Naming a statute is not teaching it, and pointing at its owner is the assertion that you did not teach it.** A section that names the Act and gives no provision, no procedure and no locator has discharged its obligation exactly. A section that adds one more paragraph has started writing a second copy of somebody else's chapter.

4.2 What legitimately remains here, and why it is not written

There is real work on this side of the line, and it is a counting job rather than an interpretive one. Statutory machinery has a running cost, and that cost is a service fact. An entitlement to make an advance directive implies somewhere it is lodged, somebody who retrieves it at the point of admission, and a record system that can produce it under pressure. An office such as the nominated representative implies people who know it exists, staff trained to identify one, and documentation that survives a transfer between institutions. Supervisory bodies imply members, sittings, quorum and a caseload. None of that is law; all of it is establishment, workload and coverage, which is what a public health chapter measures.

- Sites at which the Act's instruments can actually be lodged, retrieved and acted on, counted against the sites that admit patients.
- Staff trained to operate the machinery, counted against the establishment a district service is sanctioned for.
- Supervisory bodies constituted and sitting, counted against the units they are meant to supervise, with the reference year attached.
- The interval between an entitlement taking legal effect and the service being resourced to deliver it.

Each of those is a legitimate sentence for this chapter, and each becomes a defect the moment it is written without a source, a reference year, a population and a denominator. This chapter's evidence ledger holds no such dated, denominated quantity for the implementation load of this Act. The consequence is stated plainly rather than covered over: the argument that the Act's machinery is unevenly resourced is already in print in the legislation chapter and is not repeated here, and the quantification that would be this chapter's own contribution to it does not yet exist in a form that can be cited.

TRAP **Answering a service question with a statutory one.** Asked what the Act changed for a district service, candidates recite the rights. The examiner has already heard the rights. What distinguishes an answer is the operational consequence: who now has to be found, what now has to be recorded, and what has to be staffed for the entitlement to be real. If you cannot put a number on that, say so and say why, because at present nobody can.

Name	Point	Stop
THE ACT, AS THE LEGAL FRAME A SERVICE RUNS INSIDE	TO THE CHAPTER THAT TEACHES ITS PROVISIONS	NO PROVISION, NO PROCEDURE, NO LOCATOR HERE

MNEMONIC

Name it, point at it, price it

The three moves of a boundary section. The first two are always available. The third is the only one that belongs to this chapter, and it is the one that needs evidence this chapter does not hold for this Act.

The Mental Healthcare Act is named here as the legal frame the programmes operate inside; its provisions, instruments and procedures are taught in the Mental health legislation and forensic psychiatry notes, and this section reproduces none of them.

5 · Rights-based provisions and decriminalisation of suicide under the MHCA 2017

One provision of the Mental Healthcare Act changed the legal status of the suicide attempt in India, and it is the precondition for everything this chapter says about population suicide prevention. Before it, a person who survived an attempt met the health system through a door that also led to a police station. After it, the attempt is treated in law as the presentation of a person under severe stress rather than as an offence to be tried. That change is the reason a national prevention argument can be made at all in this country, because a population strategy cannot ask people to come forward while the coming forward is itself an offence.

The provision itself, its wording, the presumption it creates, the offence it displaces and the duty it places on government are taught in the Mental health legislation and forensic psychiatry notes, and the way the same provision operates on the individual suicide attempt is taught in the Somatic-symptom, sexual, impulse-control disorders and suicide notes. Both limbs are named here and neither is reproduced. This section is the bridge between them and the programmes that follow.

5.1 Why a public health chapter names a legal change and then stops

The rights architecture of the Act and the **decriminalisation** of the attempt are single texts with fixed wordings, and a statutory entitlement paraphrased twice in one book is an entitlement that will eventually be printed at two strengths. The risk is not abstract. The difference between a duty and a discretion lives entirely in the modal verb, and a public health author reaching for a summary sentence is precisely the author most likely to soften one. So the boundary here is drawn hard: no wording, no procedural step, no locator, no penalty, and no restatement of the duty that the appropriate government carries. Those belong to the two chapters named above and are complete there.

RULE **The fence on this construct holds wherever you approach it from.** The same provision is reached from the suicide prevention strategy, from psychiatric rehabilitation and from the emergency presentation, and it is fenced identically each time. A reader who meets it in three different sections should meet the same pointer three times, not three summaries that have drifted apart.

5.2 The service consequence, which is what remains here

Decriminalisation is a legal fact with an epidemiological shadow, and the shadow is this chapter's business. Three service consequences follow from it, and none of them requires reproducing a line of the statute.

- **Presentation.** A legal deterrent to disclosure suppresses presentation, and suppressed presentation is invisible in service data. When the deterrent is removed, recorded contacts can rise while the true rate is unchanged or falling. A service that reads a rise in attendances after a legal change as a rise in incidence has misread its own denominator.
- **Recording.** National mortality and morbidity data on suicide in India come through a police-recorded route, and what a legal category counts is not what a clinical category counts. The counting problem is taught with the national figures in the sections on suicide epidemiology and on farmer suicide in this chapter.
- **Obligation.** The provision does not merely remove a penalty; it places an affirmative duty of care on the state. A duty of that kind is discharged by services that exist, are staffed and are reachable, which is a coverage question, and coverage is what the programme sections of this chapter measure.

The third of those is where a candidate can distinguish themselves and where this chapter is most exposed. Naming the duty is the legislation chapter's work. Quantifying its discharge, meaning dated and denominated counts of the services that would actually deliver it, is this chapter's work, and it is done in the sections that own the national and district programmes rather than here.

TRAP **Answering "what did decriminalisation change" with the law rather than with the service.** The legal answer is already owned and already examined elsewhere. What a public health answer adds is the mechanism by which a legal change alters who arrives, what gets recorded and what the state is now obliged to provide. Candidates who recite the provision have answered a different question competently.

Named	Owned	Kept
THE LEGAL PRECONDITION FOR POPULATION PREVENTION	ELSEWHERE, IN BOTH ITS LEGAL LIMBS	PRESENTATION, RECORDING AND THE DUTY TO PROVIDE

MNEMONIC

Who arrives, what is written down, what is owed

The three service consequences of a legal change, and the only three this section carries. Everything else about the provision is taught in the two chapters named here.

Decriminalisation of the suicide attempt is named here as the legal precondition for population suicide prevention in India; its wording, its presumption, the offence it displaces and the duty it creates are taught in the Mental health legislation and forensic psychiatry notes and, as they operate on the attempt itself, in the Somatic-symptom, sexual, impulse-control disorders and suicide notes.

6 · Rights of Persons with Disabilities Act 2016 and mental illness as disability

Rehabilitation in India is not only a clinical activity; it is an entitlement system, and the Rights of Persons with Disabilities Act is the instrument that creates the entitlements. Everything this chapter later says about supported living, supported employment, day care, family-based care and long-term community support sits downstream of a legal architecture that recognises mental illness as a disability and attaches benefits to that recognition. This section names that architecture and points to where it is taught.

The Act itself, the conditions it recognises, the threshold at which a disability becomes a benchmark disability, the certification process, the Indian Disability Evaluation and Assessment Scale used within it, and the benefits, reservations and guardianship arrangements that follow, are taught in full in the Mental health legislation and forensic psychiatry notes. No schedule, no percentage, no procedure, no instrument domain and no score range is reproduced here. The whole band belongs there and it is complete there.

6.1 The one distinction worth carrying away from this boundary

An entitlement has two separate existences and candidates routinely collapse them. There is the **statutory entitlement**, which is what the law says a person is owed, and there is **uptake**, which is the proportion of the people who qualify who actually hold the document that unlocks it. The first is fixed the day the statute commences. The second is a service and administrative outcome that can stay near zero for years afterwards while the entitlement remains, on paper, complete.

That distinction is the reason a rights statute appears in a public health chapter at all. Legislation converts a clinical population into an eligible population. Certification converts an eligible population into a benefiting population. The gap between those two conversions is a coverage problem with all the ordinary features of one: it is unevenly distributed, it rewards the literate and the connected, it costs the applicant time and money, and it is invisible in any dataset that counts only the people who succeeded.

RULE **A right is a legal fact; a certificate is a service outcome; a benefit received is a population outcome. They are three different numbers and only the first is guaranteed by the statute.** An answer that moves from the Act straight to the benefit has skipped the step where Indian implementation actually fails.

6.2 What remains here, and the shape of the missing evidence

The legitimate remainder on this item is reach, and reach is a ratio. A defensible sentence about disability certification for mental illness in India would name an issuing authority, a

reference year, a numerator of certificates issued for mental illness, and a denominator of people estimated to meet the threshold. It would then say what proportion of the eligible population holds the document, and how that proportion varies by state, by sex and between rural and urban residence.

- Certificates issued for mental illness, nationally, in a stated year, from a named issuing authority.
- An estimate of the population meeting the threshold, derived from a survey whose own reference year and population are stated.
- The ratio of the two, presented as coverage and never as prevalence.
- The interval between application and certificate, and the proportion of applications that fail, because a process people abandon has the same population effect as a process that refuses them.

This chapter's evidence ledger holds none of those quantities for this item. That is stated here in prose rather than hidden in a footnote, because it is examinable in itself: a resident who can say that the entitlement is statutory, that the binding step is certification, and that the national uptake ratio is not established in a citable form has given a better answer than one who quotes a benchmark percentage from a slide and attaches it to the wrong denominator.

TRAP **Quoting a disability threshold as though it were a prevalence, or as though holding it were automatic.** The threshold is a legal criterion applied to an individual assessment; it says nothing about how many people cross it and nothing about how many of those who cross it are certified. The threshold, the assessment instrument and the certification procedure are all taught in the Mental health legislation and forensic psychiatry notes and none of them is reproduced in this chapter.

Named	Owned	Kept
THE ACT, AS THE ENTITLEMENT FRAME FOR REHABILITATION	SCHEDULE, THRESHOLD, INSTRUMENT AND BENEFITS, ELSEWHERE	UPTAKE AND REACH, WHICH ARE NOT YET GROUNDED

MNEMONIC

Eligible, certified, benefiting

Three shrinking populations, one statute. The law fixes the first, the certification process fixes the second, and only the third is a service outcome anybody experiences. This chapter owns the shrinkage; the statute belongs elsewhere.

The Rights of Persons with Disabilities Act is named here as the entitlement frame behind Indian rehabilitation and community services, and is taught in the Mental health legislation and forensic psychiatry notes; what belongs to this chapter is uptake, meaning certificates issued against estimated need, and no such national ratio is available to be printed.

7 · National Mental Health Policy and Tele-MANAS initiative

A policy states a direction and a programme builds something; India acquired one of each within a decade, and the pair is the cleanest worked example of the difference in this whole chapter. The National Mental Health Policy set out what the country said it wanted mental health care to become. **Tele-MANAS**, the national tele mental health programme, is a piece of national infrastructure that was actually built, that runs continuously, and that publishes reach figures. Teaching them together lets you see exactly where a stated intention turns into a countable service, and exactly where the counting starts to mislead.

Two internal boundaries apply throughout. The clinical workflow of a remote consultation and the regulatory framework governing it are taught in this chapter's section on telepsychiatry, and evidence on apps and digital products is taught in its section on digital mental health. This section owns the policy instrument and the national programme architecture, and nothing else.

7.1 The Policy: what a policy instrument is for

India's first national mental health policy was launched in **October 2014**. Being first matters more than the date itself. Until then the country had a programme without a policy, which is an odd arrangement: an implementation machine running without a published statement of what it was implementing towards. A policy supplies that statement. It names the values, the priorities, the populations regarded as needing particular attention, and the sectors expected to contribute.

What a policy cannot do is compel. It creates no right, imposes no duty and releases no money. Its influence is exerted through three indirect routes: it gives programme designers something to point at when they ask for resources, it gives advocates a published standard against which to measure performance, and it sets the vocabulary in which subsequent instruments are written. Judged as a legal instrument a policy is weak; judged as an agenda-setting instrument it can be decisive.

RULE **Ask of any policy what it would take for it to have been implemented, and whether anybody is required to report on that.** A policy with no named indicator and no reporting obligation cannot be shown to have failed, which sounds like protection and is actually the reason it can be quietly abandoned.

7.2 The scope the Policy claimed, and why that scope is unusually wide

The most consequential sentence in the Policy is not about services at all. Its stated scope includes **addressing the social determinants of mental health, including poverty, environmental issues and education**, as reproduced in the official national survey report's account of the mental health system. A policy that commits itself to social determinants has committed to something no health department can deliver alone, and that is deliberate rather than careless.

Three consequences follow, and they are examinable.

- **Intersectoral by construction.** If poverty and education are inside the scope, then departments of rural development, social welfare, labour and school education are

implementing agencies whether or not they know it. The health department becomes a coordinator as well as a provider.

- **The indicator problem becomes hard.** Service indicators are countable: contacts, admissions, prescriptions. Social determinant indicators belong to other sectors, move slowly, and cannot be attributed to a mental health policy even when they improve.
- **The accountability question sharpens.** A wide scope with no cross-sectoral reporting obligation produces a policy that is easy to agree with and impossible to audit. That is the single most useful criticism a candidate can make of any national mental health policy, and it applies well beyond India.

Which populations the Policy identifies as needing particular attention, and one population that a published policy analysis found it does not name among them, are examined in the Perinatal/women's mental health and emergency psychiatry notes. That finding is presented there as an analysis of the instrument rather than as a statement made by the instrument, and this chapter does not re-argue it.

7.3 Tele-MANAS: infrastructure rather than intention

The Government of India launched the national tele mental health programme, **Tele-MANAS**, on **10 October 2022**. Understand first what problem it is shaped around. India's specialist workforce is small and geographically concentrated, and it cannot be enlarged quickly. A telephone service breaks the geographical constraint immediately: it decouples where the caller is from where the responder is, it operates continuously rather than during clinic hours, and it can be staffed from wherever trained people happen to live.

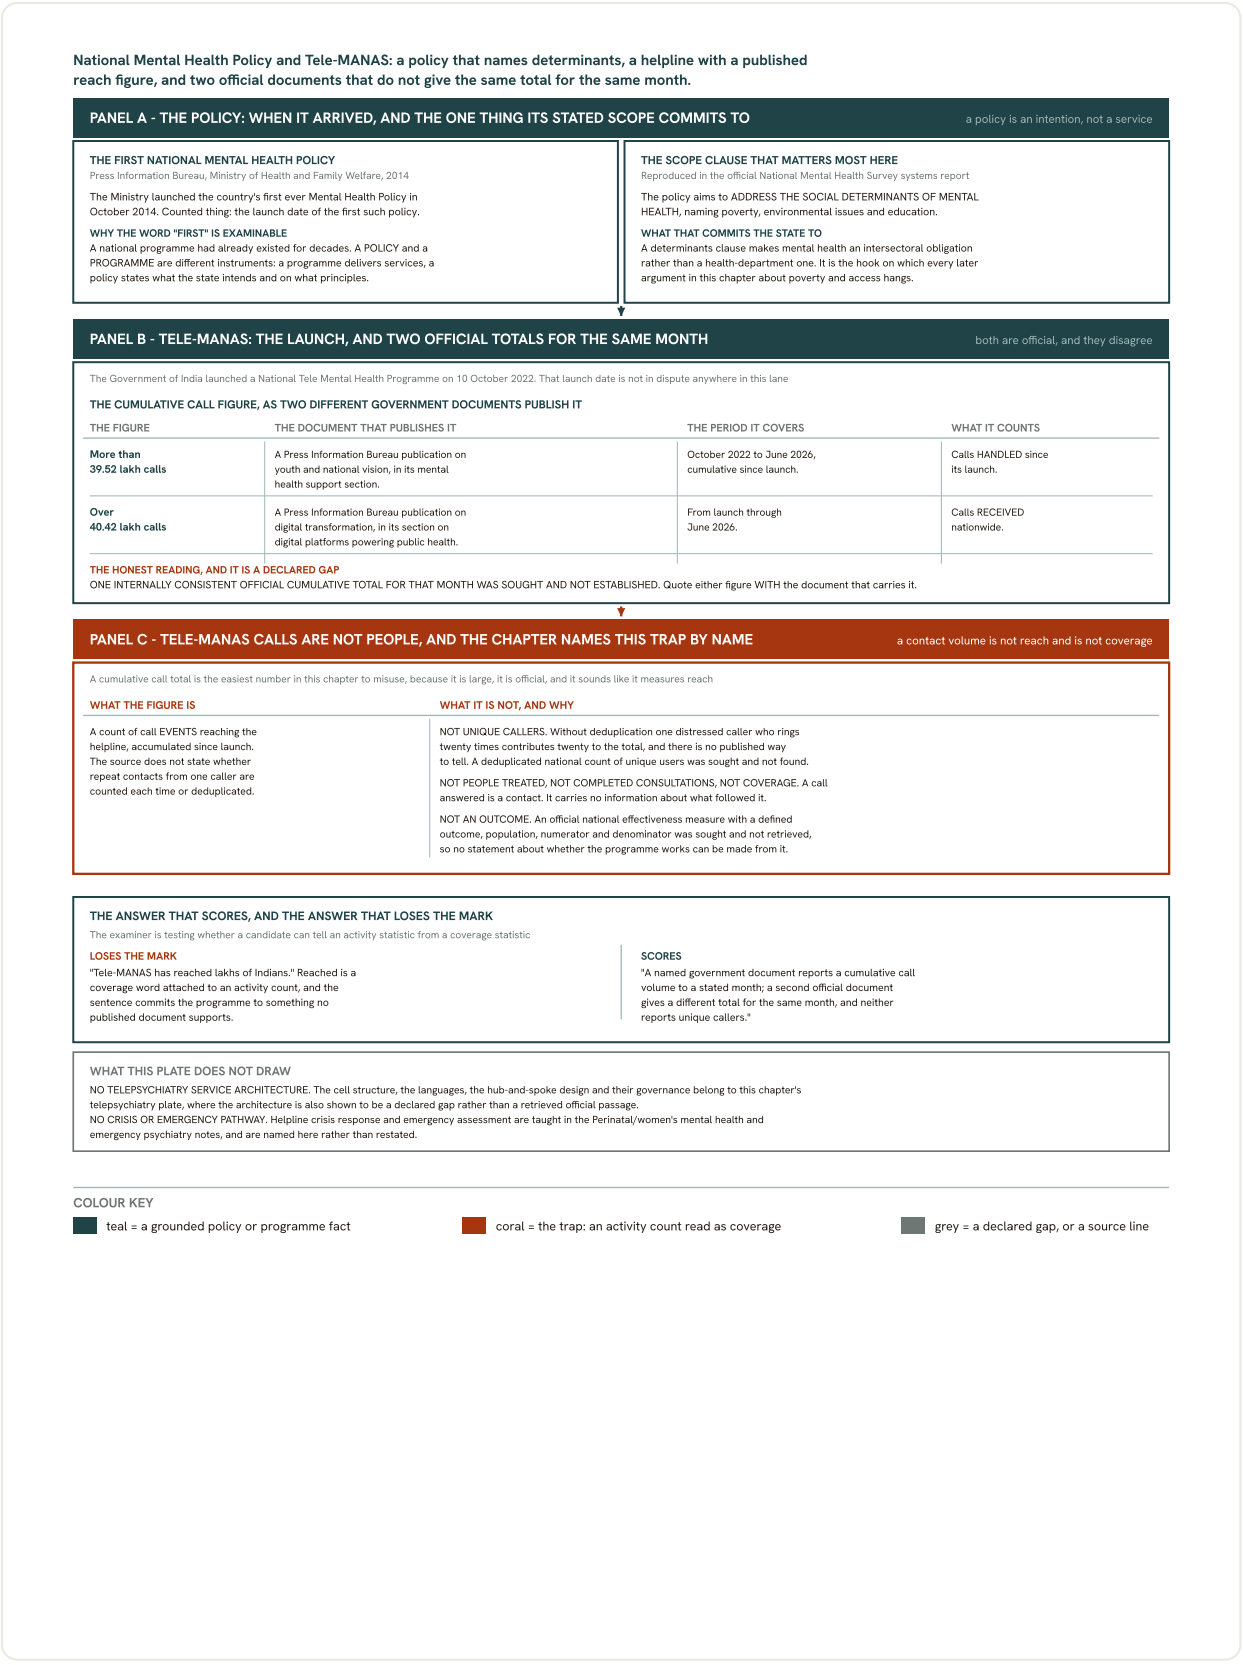


Fig 39.3 Policy and Tele-MANAS: an intention, a helpline, and a number that two official documents report differently. The country's first national Mental Health Policy was launched in October 2014 and its stated scope commits to addressing the social determinants of mental health, naming poverty, environmental issues and education. The national tele mental health programme was launched on 10 October 2022. For the same month one government publication reports more than 39.52 lakh calls handled since launch and another reports over 40.42 lakh calls received; one internally consistent official total was sought and could not be established, so either figure must be quoted with the document that carries it. Neither is a count of people: no deduplicated national count of unique users exists and no official effectiveness measure with a stated numerator and denominator was retrieved.

The accompanying figure sets out the policy instrument and the programme architecture together, including the tiered structure by which a call is answered, escalated and routed onwards, and it sets out alongside them exactly what the published reach figures do and do not count. Read the architecture as three separable claims: that a call is answered, that it is answered by somebody competent, and that it connects to onward care where onward care is needed. Only the first of the three is directly visible in any published figure.

7.4 The reach figures, and the disagreement between two official documents

This is the passage to read twice, because it is the chapter's whole discipline compressed into one construct. Two official documents report the cumulative volume this service has handled, and they do not agree.

One official release reports that Tele-MANAS handled **more than 39.52 lakh calls** by June 2026, counted cumulatively since launch, the reference period running from October 2022 to June 2026. A separate official release reports that the service had received **over 40.42 lakh calls** nationwide as of June 2026, again counted from launch. Two arms of the same official information apparatus, the same construct, the same stated month, two different totals.

10 Oct 2022	39.52 lakh	40.42 lakh
TELE-MANAS LAUNCH	CALLS HANDLED SINCE LAUNCH, ONE OFFICIAL RELEASE, TO JUNE 2026	CALLS RECEIVED SINCE LAUNCH, A SECOND OFFICIAL RELEASE, SAME MONTH

Do not reconcile them. There is no internally consistent official cumulative total for that month available to this chapter; the programme's own live dashboard was sought and could not be reached. The correct answer names both, attributes each to its own release, notes that one speaks of calls handled and the other of calls received, and says that the difference between those two verbs is not defined in either document. A resident who can do that has understood something more valuable than a memorised figure: that two official numbers for one construct usually differ because they are counting slightly different events, and that the difference is invisible unless you read the noun each one attaches to.

7.5 What a call total is not

Both figures count call events. Neither counts people. No official deduplicated national count of unique users of this service since launch was obtained for this chapter, and it was sought from the ministry and from the programme's own dashboard. That is not a small omission; it is the difference between a reach statistic and a workload statistic.

QUESTION ABOUT A NATIONAL HELPLINE	WHAT IT NEEDS	EVIDENCE POSITION IN THIS CHAPTER
How much work does it do?	A cumulative count of call events, with a reference period	Two official cumulative totals, differing, both printed above with their sources
How many people has it reached?	A deduplicated count of unique callers	Sought from the ministry and the programme dashboard; not obtained
What proportion of those who need it use it?	Unique callers over a population estimate of need	Not computable, because the numerator above does not exist
Does it help?	An operationally defined outcome, with numerator, denominator and reference period	Sought; no official national effectiveness measure was obtained

Work down that table and the shape of the evidence becomes clear. The service publishes the number that is easiest to produce and hardest to misinterpret innocently. Each successive row is more useful to a planner and less available. That pattern is not peculiar to this programme; it is what happens whenever a service is monitored by the system that runs it rather than by one that evaluates it.

TRAP **Converting calls into people.** A cumulative call total counts contacts, and neither source states whether repeat contacts from the same caller are counted each time or removed. One distressed person calling weekly for a year contributes many calls and one user. Writing that the service has "reached" that many people, or dividing the total by the population to produce a coverage rate, manufactures a statistic that no document supports. Say "calls handled, cumulative since launch, as reported by" and then name the release.

IN INDIA

When you are asked to justify a tele mental health service in a proposal or a viva, quote the volume as volume and then say what you would need to justify it as reach: a deduplicated user count, a population denominator and a defined outcome. Both published totals here are cumulative call counts, more than 39.52 lakh in one release and over 40.42 lakh in another for the same month, and neither is a count of people. Building a service case on the larger figure without that qualification is how a planning document acquires an error nobody can later trace.

7.6 Effectiveness, and the measure that is not published

The question that matters clinically is whether calling the service leaves people better off. Answering it requires an operationally defined outcome, a stated population, a numerator and a denominator, and a reference period. No official national effectiveness measure of that kind was obtained for this chapter, despite searches of the programme dashboard and ministry statements.

It is worth being precise about what that absence does and does not mean. It does not mean the service is ineffective; nothing here supports that claim either. It means the evidence that would settle the question in the direction of either answer is not published in a citable form, and that any confident statement about the programme's effect, positive or negative, is currently an

opinion. For an examination, the sequence to give is: what is published, what is not, and what would be needed. That sequence works for every programme in this chapter.

MNEMONIC

Volume, users, coverage, outcome

Four questions about any national service, in increasing order of usefulness and decreasing order of availability. Tele-MANAS publishes the first. The second, third and fourth were sought for this chapter and not obtained.

PITFALL

Treating a helpline as a substitute for a service rather than as a front door to one. A call answered competently is a triage and a beginning; its value depends almost entirely on whether the person can be connected to continuing care afterwards, which is a district-level question. A national telephone layer built over a thin district layer moves the bottleneck rather than removing it, and it makes the bottleneck harder to see, because the call is counted and the failure to connect onwards is not.

The Policy states a direction and creates no duty; the national tele mental health programme is real infrastructure whose published totals count calls, cumulatively since launch, and never people.

8 · Narcotic Drugs and Psychotropic Substances Act, Protection of Children from Sexual Offences Act and Juvenile Justice Act as they intersect psychiatry

Three statutes reach into psychiatric practice from outside psychiatry, and each one hands a mental health service a population it did not choose and a duty it did not write. The narcotic drugs legislation, the NDPS Act, defines who is an offender and who is a patient in substance use. The child sexual offences legislation, POCSO, places an affirmative reporting duty on the doctor who suspects abuse. The juvenile justice legislation creates forums that need psychiatric expertise and provides for children in institutional care who turn out to be mentally ill or dependent on substances. None of the three was written by health administrators, and all three determine what a health service has to be able to do.

This section teaches the public health and service interface only: prevention, referral, the population consequences of a legal frame, and the intersectoral machinery each statute presumes. The statutes themselves, their provisions, classifications, procedures, penalties and time frames, are taught in the Mental health legislation and forensic psychiatry notes, and the child protection procedure and the juvenile justice machinery are taught in the Child psychiatry: assessment, treatment, services and protection notes. No provision is reproduced here.

8.1 One question asked of three statutes

The productive way into this material is not to summarise three Acts, which is what an unstructured answer does and what the owning chapters already do properly. It is to ask each

statute the same two questions and compare the answers.

STATUTE	THE POPULATION IT CONSTRUCTS	WHAT IT ASKS OF A MENTAL HEALTH SERVICE
Narcotic drugs legislation	People who use controlled substances, split between an offender frame and a treatment frame	Treatment capacity that a person can reach without the reaching being risky, and services able to operate lawfully within the drug-control rules
Child sexual offences legislation	Children who have been or are being sexually abused, and the adults who encounter them	A clinician who recognises and reports, and a service organised so that reporting leads somewhere
Juvenile justice legislation	Children in conflict with law and children needing care and protection	Expert assessment capacity for its forums, and a route for children in institutional care who are mentally ill or substance dependent

Read across the third column and a single service requirement appears three times: someone competent has to be reachable, and there has to be somewhere for the person to go afterwards. That is a district service question, and it is why these three statutes belong in a public health chapter at all rather than only in a legal one.

RULE **A statutory duty on a clinician is only as real as the service standing behind it.** Legislation can require a doctor to act; it cannot by itself create the trained person, the pathway or the receiving service that makes the action useful. The gap between the duty and the capacity is a health system problem, and it is the part of this material that belongs here.

8.2 Drug control as a public health instrument

Every drug-control regime does two things at once, and the two pull in opposite directions for a health service. Supply reduction attacks availability through prohibition, interdiction and penalty. Demand reduction attacks consumption through prevention, treatment and rehabilitation. A statute that carries both creates a population that is simultaneously a target of enforcement and a target of care, and the person in that population has to decide which of the two they are approaching when they walk into a clinic.

That ambiguity has measurable service consequences and they are the substance of this subtopic.

- **Delayed presentation.** Where disclosure carries legal risk, people present later, sicker, and often through an emergency rather than through a clinic. The delay is invisible in service data because the earlier presentations that never happened leave no trace.
- **Incomplete disclosure.** A person who fears legal consequence under-reports quantity, route and the other substances used. Every clinical decision downstream, including dose, is then made on a corrupted history.

- **Family and employer concealment.** Legal risk attaches to the household as well as the individual, so the informants a clinician relies on have their own reason to minimise.
- **Distorted surveillance.** Enforcement data count seizures and cases; health data count contacts. Neither counts prevalence, and a rise in either can reflect activity rather than epidemiology.

The reformative side of the drug-control frame, meaning the statutory routes by which a person who undergoes treatment is treated differently by the law, and the arrangements for recognising treatment institutions, are taught in the Mental health legislation and forensic psychiatry notes. What that chapter also already establishes is that such routes work only where a qualifying institution is genuinely within reach, so this section does not re-derive that point and instead takes it as given.

8.3 Where the treatment actually has to sit

If reach is the binding constraint, then the design question is what level of the health system can lawfully and safely provide substance use treatment. National operational guidance for mental, neurological and substance use care at the first level of care takes a clear position on the hardest case. It states that **opioid substitution therapy may be provided by adequately trained medical officers at primary health centres and urban primary health centres, in areas with high levels of opioid dependence, in keeping with the drug control Act and Rules**.

Read the qualifiers, because they carry the whole design. The provider is a medical officer, not a specialist, which is task-shifting applied to the most tightly regulated treatment in psychiatry. The condition is adequate training, which converts the sentence from a permission into a training obligation. The targeting is geographical, directed at areas where opioid dependence is concentrated rather than spread thinly everywhere. And the whole permission sits inside the drug-control rules, so the regulatory frame is not bypassed but worked within.

> **IN INDIA**
>
> Where opioid dependence is concentrated, do not design the service around a specialist centre that people travel to. National operational guidance places this treatment at the first level of care, with adequately trained medical officers at primary health centres and urban primary health centres as the providers, within the drug-control rules. The practical decisions that follow are about training, supervision and regulated supply at that level, not about building another specialist clinic; a treatment that requires daily or near-daily attendance fails on distance long before it fails on efficacy.

8.4 Mandatory reporting as a service obligation, not only a personal one

The child sexual offences legislation is the point at which a clinician's private judgement stops being private. Official guidance issued under that legislation is explicit about the clinician's position: **a doctor who has reason to suspect that a child has been or is being sexually abused is required to report it**. Two features of that sentence do the work. The trigger is suspicion with a reason behind it, not certainty and not proof, which is a deliberately low

threshold. And the verb is one of requirement, not of discretion, so the decision the clinician faces is not whether to report but how to do it well.

The procedural detail, meaning who the report goes to, what form it takes, what protection the reporter has and what follows for the child, is taught in the Child psychiatry: assessment, treatment, services and protection notes. What belongs here is the systems consequence, which is routinely underestimated.

- **The duty falls on every clinician, so the service must assume every clinician will meet it.** A duty discharged well by the one consultant who happens to know the pathway is a service that has not implemented anything.
- **Confidentiality is reshaped rather than suspended.** The child and the accompanying adult need to be told, in advance and in plain language, what a clinician cannot keep to themselves. Saying it before the disclosure is honest; saying it afterwards is a betrayal the family will remember.
- **Reporting is the beginning of a service obligation, not the end of one.** The child still needs assessment, treatment and follow-up, and the reporting event frequently disrupts exactly the therapeutic relationship in which that would have happened.
- **A service needs a named route and a rehearsed one.** Which staff member acts, whom they contact, what is recorded and who supports the family afterwards should be settled before the first case, because the first case will present out of hours to the least experienced person present.

TRAP **Treating mandatory reporting as an ethical dilemma to be weighed case by case.** Candidates reach for the confidentiality argument because it is the interesting one. Where an official duty to report exists, the weighing has already been done by the legislature and the clinician's judgement is engaged in how to report, how to prepare the child and family, and how to preserve continuity of care afterwards. The examinable skill is the service design, not the deliberation.

8.5 Juvenile justice: two places psychiatry is written into the system

The juvenile justice framework creates forums that a mental health service interacts with directly. The Juvenile Justice Board and the Child Welfare Committee are, from a service point of view, destinations: places a child is routed to and routed from, each with its own decisions to make and its own need for information a clinician can supply. Their constitution, composition and procedure are taught in the Child psychiatry: assessment, treatment, services and protection notes.

Two provisions of that legislation reach psychiatry directly and both are service-shaping rather than merely procedural.

The first concerns expertise. In carrying out its preliminary assessment, **the Board may obtain assistance from experienced psychologists, psycho-social workers or other experts**. The modal is permissive and that is the important part. The statute contemplates expert input rather than compelling it, which means whether the expertise is actually available in a district is a

service question and not a legal one. A permissive provision in a district with no available expert is a provision that does nothing.

The second concerns children already inside the system. The legislation expressly provides for **the transfer of a child who is mentally ill or addicted to alcohol or other drugs**. That is a statutory recognition that institutional care will contain children who need psychiatric or addiction treatment, and it presumes a receiving service exists to transfer them to. Whether it does is, again, a district capacity question.

One further limb of the same legislation, the adoption route and the psychiatric fitness of prospective adoptive parents, is named here in a single clause because a reader tracing the statute will look for it, and it is taught in the Mental health legislation and forensic psychiatry notes. No eligibility criterion, regulation or fitness assessment procedure appears in this chapter.

Reach	Report	Receive
DRUG CONTROL: TREATMENT A PERSON CAN SAFELY APPROACH	CHILD PROTECTION: A DUTY ON EVERY CLINICIAN	JUVENILE JUSTICE: EXPERTISE AND A PLACE TO TRANSFER TO

8.6 What the three have in common, and what a service can do about it

All three statutes make a mental health service a downstream consumer of categories it did not define. The categories are legal, the thresholds are legal, and the timing is set by legal processes. A service that treats this as somebody else's problem discovers its consequences one crisis at a time. A service that plans for it does four things, and none requires a change in the law.

- **Identify the interfaces in advance** and name, for each, the person in the service who acts and the external counterpart they contact.
- **Train for the trigger rather than the statute.** Staff need to recognise the situation that starts the duty; the legal detail can be looked up, and the recognition cannot.
- **Build the receiving capacity, not only the referring pathway.** Every one of these statutes presumes somewhere to send a person. Pathways into services that do not exist generate paperwork and nothing else.
- **Record at the interface.** Cross-sector work is where documentation is weakest and where it is most likely to be examined later.

MNEMONIC

Reach, report, receive

What each of the three statutes asks a mental health service for: treatment a person can safely approach, a clinician who recognises and reports, and a service able to receive a child the system sends. Every one of the three is a capacity question rather than a legal one.

PITFALL

Assuming an intersectoral pathway exists because a statute presumes one. Legislation routinely presumes a counterpart service, a trained expert or a receiving institution, and says nothing about creating any of them. Before promising a court, a committee or a family that a route exists, confirm that the person at the other end of it is currently in post. The commonest failure at these interfaces is not a wrong decision; it is a correct referral into a vacancy.

These three statutes are named here and taught elsewhere; what this chapter owns is the capacity each one presumes, which is treatment a person can safely reach, a clinician who recognises and reports, and a service able to receive the child the system sends.

9 · Ayushman Bharat, insurance parity for mental illness and the IRDAI mandate

Who pays decides who is treated, and in mental health the answer has historically been the household. A person with a chronic psychiatric illness in India typically needs medication continuously for years, review appointments indefinitely, occasional admission, and travel to reach any of it, while earning less because of the illness. Every one of those costs falls on a family whose earning capacity the illness has already reduced. Financing is therefore not an administrative appendix to this chapter; it is one of the determinants of whether care is used at all.

This section teaches the financing and insurance machinery: how publicly financed health insurance is built, what parity means and how it is implemented, why mental illness fits an insurance model badly, and what can and cannot currently be said with evidence about coverage and utilisation in India. The legal right to **insurance parity**, its statutory wording and the account of how that right came in practice to move the insurance regulator, are taught in the Mental health legislation and forensic psychiatry notes. What that chapter does not name is the regulator, and this one does: the Insurance Regulatory and Development Authority of India, the **IRDAI**.

One boundary inside the Ayushman Bharat umbrella needs stating at the outset, because the shared brand causes real confusion. This section teaches the insurance and benefit-design limb, meaning the publicly financed hospitalisation scheme and the regulation of insurance products. The school health and wellness limb of the same umbrella is a different programme with a different delivery mechanism, and it is taught in the Child psychiatry: assessment, treatment, services and protection notes. One brand, two programmes, and their figures must never be pooled.

9.1 Four ways a health system pays, and what each does to psychiatric care

Any health system pays for care through some combination of four mechanisms, and each has a characteristic effect on a chronic, stigmatised, outpatient-dominant illness.

MECHANISM	HOW IT WORKS	EFFECT ON CHRONIC PSYCHIATRIC CARE
Tax-financed direct provision	Government funds and runs facilities; care is free or nominal at the point of use	Best fit in principle, because it can cover continuing outpatient care; limited in practice by staffing and reach
Publicly financed insurance	Government buys defined packages of care for defined beneficiaries from empanelled providers	Fits admissions and procedures; fits continuing outpatient treatment poorly unless deliberately designed to
Private voluntary insurance	Individuals or employers buy a product from an insurer, subject to regulation	Historically excluded or restricted psychiatric cover; the reason a parity requirement was needed at all
Out-of-pocket payment	The household pays at the point of use	Default when the other three fail; the mechanism through which illness produces impoverishment

Read the fourth row as the residual. Out-of-pocket payment is not a policy choice anybody makes; it is what happens when the first three do not cover something. Every argument for parity, for benefit-package design and for public financing in mental health is ultimately an argument about shrinking that fourth row.

RULE **An insurance benefit is defined by the event it pays for, and psychiatric care is mostly not an event.** Hospitalisation insurance pays for admissions. Chronic psychiatric care is a long sequence of inexpensive contacts with occasional expensive ones. A benefit designed around discrete admissions can be generous, well administered and still leave the ordinary course of a mental illness entirely unfunded.

9.2 The machinery of a publicly financed insurance scheme

To criticise a scheme intelligently you have to know where its joints are. A publicly financed hospitalisation scheme has six moving parts, and a mental health service can fail at any of them.

- **Beneficiary identification.** Somebody has to be established as entitled before care, which requires a document, a database entry and usually literacy or an intermediary. People who are homeless, wandering, undocumented or estranged from family fail here first, and that population overlaps heavily with severe mental illness.
- **Empanelment.** Only listed facilities can provide the benefit. If the psychiatric facilities in a district are not empanelled, the entitlement is real and unusable.
- **Benefit package design.** The scheme pays for named packages. What is not named is not paid for, and naming psychiatric care in package terms is genuinely hard because much of it is time rather than procedure.
- **Pre-authorisation.** Approval before treatment suits planned surgery and suits acute psychiatric admission badly, since the admission decision is often urgent and made out of

hours.

- **Package rates.** A fixed price per package determines whether providing the care is financially survivable for the facility. A rate set below cost produces empanelled facilities that quietly decline the cases.
- **Claims and audit.** Payment follows documentation. Psychiatric documentation is narrative rather than numerical, which makes claims slower to settle and easier to reject.

Notice that four of the six failure points are invisible to a patient. A family that is told the scheme does not cover this is rarely told which joint failed, and the distinction matters enormously for anybody trying to fix it.

9.3 Parity, and its three separable limbs

Parity is a word that hides three different demands, and an answer that treats it as one has lost most of the available marks.

- **Coverage parity.** Mental illness is included at all, rather than excluded by a clause. This is the crudest limb and the easiest to verify.
- **Benefit parity.** Where it is included, the limits, sub-limits, waiting periods, co-payments and caps are no less favourable than those applied to physical illness. A product can satisfy coverage parity and fail here completely, and this is where formal compliance most often stops.
- **Administrative parity.** The pathway to claiming is no harder: the same documentation burden, the same pre-authorisation friction, the same likelihood of settlement. This limb is almost never measured and is where a formally compliant product can still deliver very little.

The value of separating them is diagnostic. When you are told that a market has achieved parity, the follow-up questions are which limb, measured how, and by whom. A regulator can require the first with a stroke of a pen; the second requires product-level scrutiny; the third requires claims data that nobody publishes.

TRAP **Treating a regulatory requirement as evidence that products changed.** A direction to insurers is an input. Whether the products in the market now carry psychiatric cover on comparable terms, and whether people actually claim under it, are two further questions that require product analysis and claims data. Candidates routinely answer the second and third by asserting the first more loudly.

9.4 The regulator's instrument, which this chapter could not open

Here the honest position has to be stated in full, because this section's central document is one this chapter does not hold. The insurance regulator is the body that turns a statutory entitlement into a requirement on insurance products, and it does so through circulars addressed to insurers. The specific circular concerning mental illness coverage, its reference number, its date and its verbatim operative requirement, was sought for this chapter through the regulator's own health department circular index, through its master circular on health insurance business, and

through searches for the specific circular references that circulate in secondary discussion. It was not obtained.

What follows from that is a set of refusals, and they are deliberate. This chapter prints no circular number, no circular date and no quotation of what the regulator directed. It does not paraphrase the requirement, because the strength of a regulatory obligation lives in its modal verb and a paraphrase of an unread document invents that verb. It does not state when the requirement took effect. And it does not assert that any particular class of product must now behave in any particular way.

The corresponding examination answer is available and is a good one. A resident can say that the statutory right sits in the mental health legislation, that the regulator responsible for translating it into product requirements is the Insurance Regulatory and Development Authority of India, that this translation is effected through circulars to insurers, and that they would cite the specific circular by number and date rather than from memory. That answer is complete, correctly attributed and contains nothing invented.

9.5 Utilisation, and the rate that does not exist

Suppose the requirement is in force and products comply. The question that then matters for a population is how much psychiatric care is actually being funded through insurance. Answering it needs four things: a numerator of claims or of people treated, a denominator of eligible insured lives, a reference period, and the resulting rate.

No such published utilisation figure was obtained for this chapter. Searches of the regulator's circular index, its health department material and its master circular on health insurance business did not yield a numerator, a denominator or a rate for mental illness coverage. The consequence is that no proportion, no growth trend and no comparison with physical illness utilisation appears anywhere in this section.

The same applies on the public side. A publication of the national authority that administers the publicly financed scheme, specifying which mental health services are covered and, if published, their national utilisation, was sought and not obtained. So this chapter names the scheme and the machinery, and does not list its psychiatric packages or quote their uptake.

Sought	Sought	Sought	None held
THE REGULATOR'S CIRCULAR, ITS NUMBER, DATE AND WORDING	A UTILISATION NUMERATOR, DENOMINATOR AND RATE	THE COVERED-SERVICES LIST FOR THE PUBLIC SCHEME	SO NO FIGURE IS PRINTED IN THIS SECTION

Four tiles, three searches, no figures. That is an uncomfortable result for a section about financing and it is reported rather than concealed, because the alternative is worse. Financing figures are the easiest numbers in this whole chapter to half-remember from a seminar slide, and a coverage percentage attached to the wrong denominator is a fabrication that no digit check will catch.

9.6 What a resident should actually be able to say

The absence of figures does not make this material unanswerable; it changes what a good answer consists of. Structure it in four moves.

- **Name the layers.** A statutory right to parity, a regulator that translates it into product requirements, a publicly financed hospitalisation scheme for a defined beneficiary population, and out-of-pocket payment as the residual.
- **Say where mental illness fits badly and why.** Outpatient-dominant, longitudinal, time-based rather than procedure-based, and expensive in aggregate through frequency rather than through unit cost.
- **Separate the three limbs of parity** and say which of them any given claim about compliance actually addresses.
- **State the evidence position honestly.** That the operative regulatory instrument and the utilisation data were sought from the regulator and the national authority and are not available to you in citable form, and that you would not quote a coverage or utilisation figure you cannot attribute.

MNEMONIC

Identify, empanel, package, authorise, price, pay

The six joints of a publicly financed insurance scheme, in the order a case passes through them. Ask which joint failed before concluding that a scheme does not cover psychiatry.

PITFALL

Reassuring a family that a scheme will cover an admission before checking the two joints that most often fail: whether the facility is empanelled for this care, and whether the beneficiary record is in order today rather than in principle. Both fail silently and both fail at the point of admission, when the family has already committed. Check them at the outpatient visit, not at the door of the ward.

Parity is three separate demands, coverage, benefit and administration; a health system that has met the first has usually not been tested on the other two, and in India the instruments that would test them were sought for this chapter and are not available in citable form.

Epidemiology, burden and treatment gap

10 · National Mental Health Survey of India: prevalence figures and key findings

Every service argument in Indian psychiatry eventually rests on one survey, and most people who quote it cannot say what it counted. The **National Mental Health Survey** is the country's reference estimate of how much mental illness there is, who has it, and how much of it is going untreated. Its figures are the denominators behind almost every planning claim in this chapter. They are also the figures most often quoted at the wrong value, for the wrong population, without the year they refer to, and against a denominator that belongs to something else.

This section teaches the survey as a planning instrument: what it sampled, how it defined a case, what it found, and what those findings can and cannot be used to decide. How the estimates were

derived and what they mean epidemiologically are taught in the Research Methods, Biostatistics and Epidemiology notes. The treatment gap as a construct, its determinants and what closes it are taught in their own section of this chapter, and the burden of mental disorders expressed in disability-adjusted life years is taught in another; no figure from either is reproduced here.

10.1 What a national survey is for, and what it is not

A national survey produces a population estimate. A service produces a service statistic. The two answer different questions and are constantly confused, so fix the difference before anything else. A service statistic tells you about the people who arrived: how many were seen, what they had, what was done. It is silent about everyone who did not arrive, and in mental health that silent group is the larger one. A population survey goes out and asks, so it sees people who never present at all. This is why a national survey is the only instrument that can tell you the size of the problem rather than the size of the workload.

What it cannot do is tell you about a district. Survey estimates are constructed for the populations the design was powered to describe. Applying a national figure to a particular district, or a state figure to a taluka, imports an assumption about similarity that the survey never tested. This is the single commonest analytical error made with these numbers in Indian service planning.

RULE **A prevalence figure is a statement about a defined population in a defined period, measured with a defined instrument under a defined case definition.** Change any one of those four and the number changes. Quote a prevalence without them and you have quoted a rumour that happens to be accurate.

10.2 The design, and why the state list matters

The survey was undertaken during **2015-16**. Its fieldwork reached **39,532 individuals**, sampled across 720 clusters from 80 talukas in 43 districts of the 12 selected states. It covered 12 states, namely Punjab, Uttar Pradesh, Tamil Nadu, Kerala, Jharkhand, West Bengal, Rajasthan, Gujarat, Madhya Pradesh, Chhattisgarh, Assam and Manipur, drawn to represent the regions of the country.

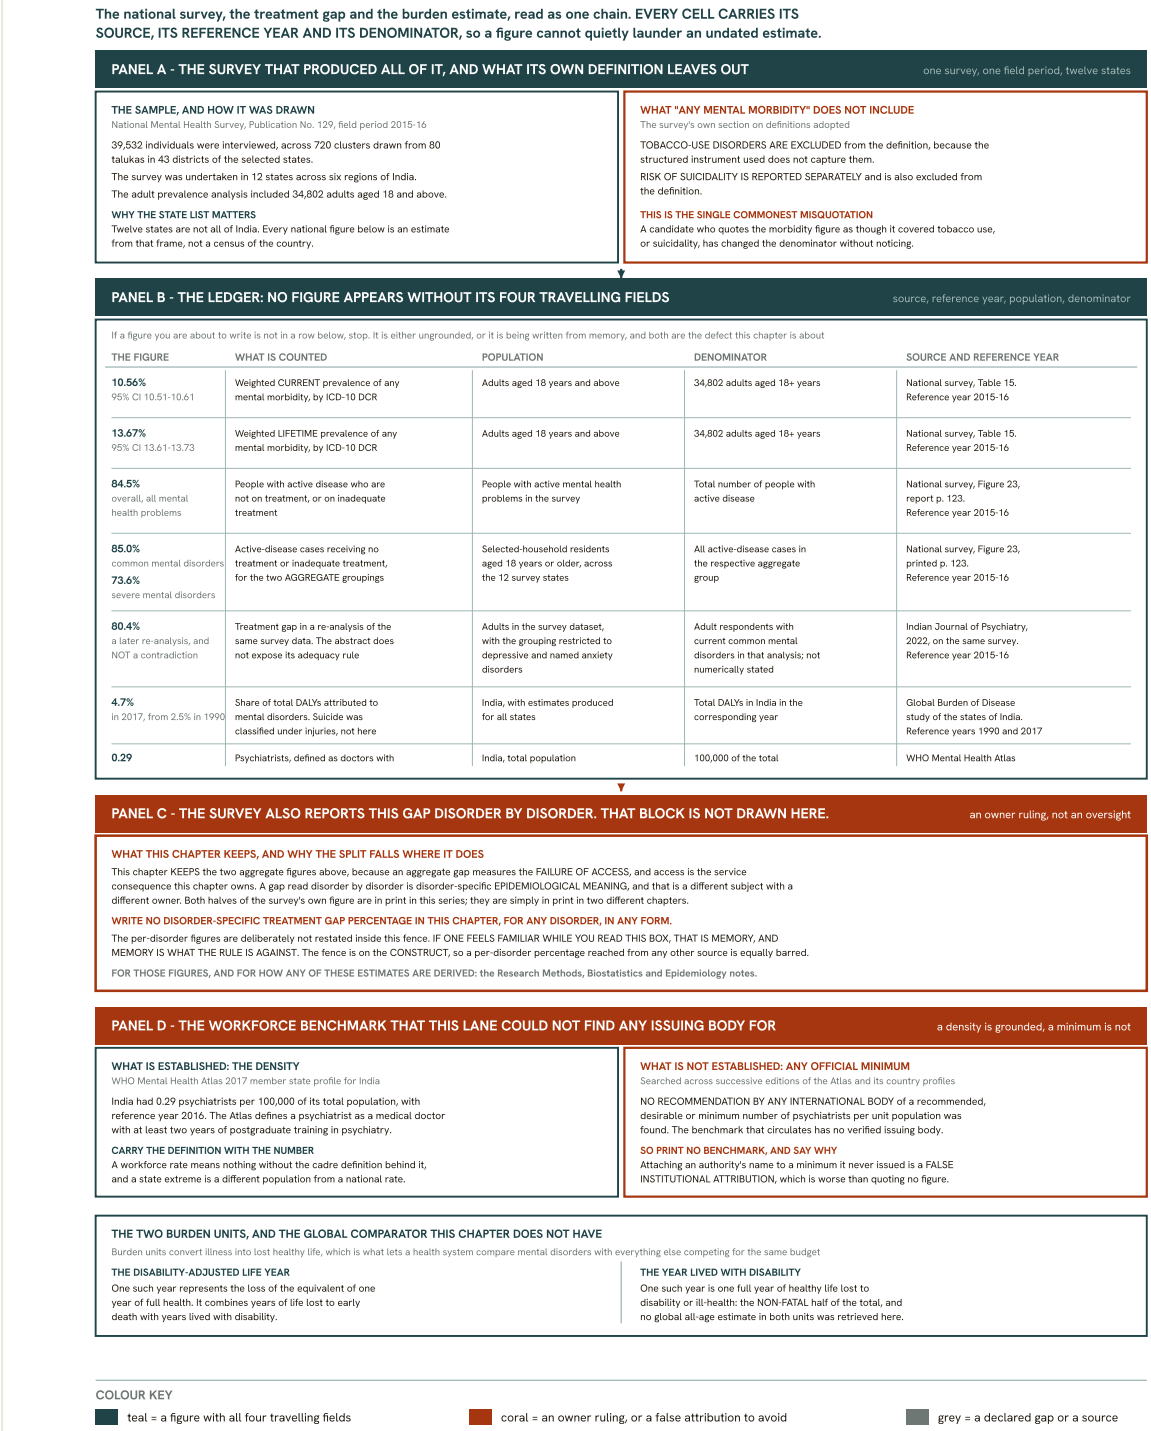


Fig 39.4 **The survey, the gap and the burden, with every figure carrying its four travelling fields.** The national survey interviewed 39,532 individuals across 720 clusters in 43 districts of 12 states, and its morbidity definition excludes tobacco-use disorders and reports risk of suicidality separately. Current prevalence was 10.56 per cent and lifetime 13.67 per cent. The treatment gap is printed here only in its AGGREGATE forms: 84.5 per cent overall, 85.0 per cent for common and 73.6 per cent for severe mental disorders, with 80.4 per cent from a later re-analysis whose disorder grouping was narrower. The survey's disorder-by-disorder block is owned by the Research Methods, Biostatistics and Epidemiology notes and no per-disorder percentage appears on this plate. Mental disorders accounted for 4.7 per cent of total DALYs in India in 2017, against 2.5 per cent in 1990. India had 0.29 psychiatrists per 100,000 population, and no international body's recommended minimum could be found, so none is printed.

The accompanying figure sets out the chain these findings sit in, from the survey's prevalence estimates through the treatment gaps derived from them to the burden expressed as health loss, with the source, reference year and denominator printed against every cell so that no figure in the chain can be read without them.

Read the state list rather than skipping it, because it is the survey's most important limitation and its most quotable one. Twelve states were sampled out of a much larger number, deliberately spread across the regions of the country so that the whole could be described from the parts. That is a legitimate and standard design. It also means that a state not on that list has no estimate of its own from this survey, and that any national figure derived from it is a construction from twelve sampled populations rather than a measurement of every state.

TRAP **Quoting the reference year as the year you are reading it.** These are figures from fieldwork conducted in 2015-16, and the reference year travels with them permanently. An estimate does not become current by being repeated, by being reprinted in a later document, or by being cited in a recent policy paper. Every time you quote one of these numbers, say the period it describes.

10.3 The case definition, which is where most misquotation begins

If you learn one thing from this section that other candidates will not know, make it this. The survey's headline category, any mental morbidity, is a constructed definition with explicit inclusions and, more importantly, explicit exclusions. Diagnoses were made against ICD-10 criteria and captured through a structured diagnostic interview, the Mini International Neuropsychiatric Interview, which fixes both what could be detected and what could not.

Two exclusions in particular change what the headline number means. **Tobacco-use disorders were excluded from the definition, and risk of suicidality, although captured, was reported separately and also excluded from it**. Both decisions are defensible and both are consequential. Tobacco dependence is extremely common, so including it would have raised the headline figure substantially and would have made it a different sort of number. Risk of suicidality is a risk state rather than a disorder, so folding it into a morbidity total would have mixed two kinds of thing.

The consequence for a reader is precise. When somebody quotes the national prevalence of mental disorders in India, the correct follow-up question is which construct they mean, because this headline excludes the commonest substance-use problem in the country and excludes the risk state that most service planning around suicide is directed at.

IN INDIA

Do not size a tobacco cessation component or a suicide prevention component from this survey's headline morbidity figure. Tobacco-use disorders are excluded from the definition and risk of suicidality is reported separately and excluded from it, so the headline number contains neither. A district plan that budgets for either on the basis of the overall morbidity estimate has sized it from a denominator that deliberately leaves it out, and will under-provide in exactly the two areas where demand is highest.

10.4 The headline findings, with their denominator attached

The adult prevalence analysis was conducted on **34,802 adults** aged 18 years and above. That figure is the denominator of the prevalence estimates, and it is smaller than the total number of people interviewed because the adult analysis is a subset of the whole sample. Carry it with the percentages; a proportion without its denominator cannot be interrogated.

Among adults aged 18 years and above, the weighted current prevalence of any mental morbidity was 10.56 per cent, with a 95 per cent confidence interval of 10.51 to 10.61. The corresponding weighted lifetime prevalence of any mental morbidity was 13.67 per cent, with a 95 per cent confidence interval of 13.61 to 13.73. Both are weighted percentages, which is stated in the report's own table footnote and is not a detail: weighting is what allows a sample drawn unevenly across strata to speak for the population it was drawn from.

The survey also reports how much of that morbidity was going untreated. Its treatment gap for mental health problems was **84.5 per cent**. Read the denominator carefully, because this is where the number is most often mangled: it is a proportion of the total number of people with active disease, counting people with active disease who are not on treatment or on inadequate treatment. It is not a proportion of the population. The treatment gap as a construct, its determinants and the interventions that narrow it are taught in their own section of this chapter, and the gap read disorder by disorder, with its epidemiological meaning, is taught in the Research Methods, Biostatistics and Epidemiology notes.

39,532	10.56%	13.67%	84.5%
INDIVIDUALS INTERVIEWED ACROSS THE SELECTED STATES	CURRENT ANY MENTAL MORBIDITY, ADULTS AGED EIGHTEEN AND OVER	LIFETIME ANY MENTAL MORBIDITY, SAME ADULT SAMPLE	TREATMENT GAP, OF PEOPLE WITH ACTIVE DISEASE

10.5 Four numbers, four denominators

Set the headline figures side by side with what each is a proportion of, and the difference between them stops being memorisable trivia and becomes a working tool.

FIGURE	POPULATION IT DESCRIBES	WHAT IT IS A PROPORTION OF	WHAT IT DOES NOT TELL YOU
Current prevalence 10.56 per cent	Adults aged 18 years and above in the sampled states	34,802 adults in the analysis	How many will develop illness later, or how many are in treatment
Lifetime prevalence 13.67 per cent	The same adult sample	34,802 adults in the analysis	How many are unwell now, which is the number a service needs
Treatment gap 84.5 per cent	People with active mental health problems in the survey	The total number of people with active disease	Why they are untreated, which is a separate section of this chapter
39,532 individuals interviewed	The whole sampled population, adults and others	Not a proportion; a sample size	Anything about states outside the sampling frame

The second row repays attention. Lifetime prevalence is always the larger and more quotable figure, and it is the less useful of the two for planning a service, because a service treats people who are unwell now. Where a planning document quotes a lifetime figure to justify current establishment, it has quietly inflated its case.

10.6 Children, and the limit this survey states about itself

The adult findings above are exactly that: adult findings, from an adult analysis, with an adult denominator. The survey carried an adolescent component, but it was a feasibility sub-sample confined to some of the survey states rather than a nationally representative adolescent survey, and it does not support an all-India childhood estimate.

The adolescent figure, the limits attached to it and the position that no established all-India psychiatric prevalence can be quoted for younger children are taught in the Child psychiatry: assessment, treatment, services and protection notes. No child or adolescent prevalence figure appears in this chapter, and no adult figure here is extended downwards to children. Extrapolating an adult rate to a child population is not a conservative approximation; it is a different claim about a different population made without evidence.

10.7 Using the findings for priority-setting without misusing them

The reason to know the design so well is that priority-setting arguments are won and lost on the denominator. Four disciplines make the difference between a survey-based argument that holds and one that collapses under a single question.

- **Name the construct before the number.** Any mental morbidity, current, adults, with tobacco-use disorders excluded and suicidality reported separately, is a specific thing. Saying "the prevalence of mental illness in India" and then giving a figure omits every part that makes it true.
- **State the reference period every time.** The estimates describe fieldwork in a stated period and nothing since. If you are asked how much of this is still true, the honest answer is that the survey cannot say, because a survey speaks only about when it was done.
- **Keep current and lifetime apart.** Use the current figure for establishment and workload arguments and the lifetime figure for arguments about population exposure over a life course.
- **Never mix denominators.** A prevalence is a proportion of a population, a treatment gap is a proportion of the people with the condition, and no arithmetic combining the two directly produces anything meaningful.

MNEMONIC

Who, when, how defined, out of what

The four fields that make a survey number a fact. Which population, which reference period, which case definition and instrument, and which denominator. Any of the four missing, and the figure cannot be checked by anybody who hears it.

PITFALL

Presenting the national estimate to a district administration as though it described their district. The survey sampled a defined set of states through a clustered design, and the national figure is built from those samples. A district that differs in its age structure, its urban proportion, its migration pattern or its substance-use profile will differ in its true prevalence, and nobody knows in which direction. Use the national estimate to argue that a service is needed, and local data, however imperfect, to argue how much.

Current weighted prevalence of any mental morbidity in adults was 10.56 per cent and lifetime was 13.67 per cent, both from fieldwork in 2015-16, on a definition that excludes tobacco-use disorders and reports suicidality separately.

11 · Treatment gap for mental disorders in India and its determinants

The treatment gap is the single most quoted number in Indian community psychiatry and the one most often quoted against the wrong denominator. It is not a measure of how much illness there is. It is a measure of how completely a health system is failing to reach the people who already have it. That distinction is the whole of this section: everything that follows is about what makes the gap the size it is, who is left on the wrong side of it, and what has actually been shown to move it.

How the estimates are derived, and what the gap means read disorder by disorder, are taught in the Research Methods, Biostatistics and Epidemiology notes. The national survey that produced the Indian estimates, its design and its own overall reported figure are taught in the section of this chapter that owns the survey. What belongs here is the causal and service account: determinants, inequities, consequences and interventions.

11.1 What the gap measures, and out of what

A **treatment gap** is the proportion of people with a condition who are not receiving treatment, or who are receiving treatment inadequate to their condition. Three features of that definition do all the work and each is a route to error.

- **The denominator is the diseased, not the population.** A gap of four fifths means four of every five people who have the condition are untreated. It says nothing about how common the condition is. Multiplying a gap by a population is arithmetic on two unrelated things.
- **Inadequate treatment counts as gap.** The measure is not simply contact. A person prescribed a subtherapeutic dose, seen once and never followed up, or treated by someone unqualified to treat them, sits inside the gap despite having reached a service. This is why closing a gap is harder than opening clinics.
- **The gap is population-specific.** It differs by disorder, by severity, by state, by sex and by wealth, so a single national number is an average across all of those and describes nobody exactly.

RULE **Never quote a treatment gap without naming the population it belongs to.** The commonest defect in this whole area is a figure carried across from the population it was calculated for to a different one. A correctly transcribed percentage attached to the wrong group is still a false statement, and it is one that no arithmetic check will catch.

11.2 The national figures, and the values that do not agree

India's national survey reports the gap for two aggregate groupings. **The treatment gap was 85.0 per cent for common mental disorders and 73.6 per cent for severe mental disorders**. The population is selected-household residents aged 18 years or older across the 12 survey states, the reference period is 2015-16, and the denominator in each case is all active-disease cases in the respective common-mental-disorder or severe-mental-disorder group. What is being counted is active-disease cases receiving no treatment or inadequate treatment.

Notice what the pair tells you before you go any further. The gap is *larger* for the common disorders than for the severe ones. That inversion surprises people who expect severity to track access, and it is one of the most instructive facts in this chapter. Severe illness is disruptive, visible to families and hard to ignore, so it eventually generates a contact even where services are poor. Depression and anxiety are survivable in a way that lets a person and a household absorb them indefinitely without anyone deciding that this is illness.

A lower figure circulates widely for the same construct, and a resident who meets it should be able to place it. A later analysis based on the same national survey data reported a treatment gap of **80.4 per cent** for its own analysis of common mental disorders. It is not a competing measurement of the same thing: that analysis restricted common mental disorders to depressive disorders and six named anxiety disorders, so its denominator is adult respondents with current common mental disorders included in that analysis and not the survey's own broader grouping. The opened abstract reports the treatment gap but does not expose its detailed numerator, its treatment-adequacy rule or its treatment window, and does not state the numeric denominator.

A third figure for common mental disorders also circulates and is not printed here. It was sought directly, in the published analysis and through several routes to its own tables, and it could not be established whether it is an overall pooled estimate or a subgroup estimate. A percentage whose population cannot be determined is not a fact yet, and this chapter does not print it in either reading.

TRAP **Treating two treatment gap figures from one survey as a contradiction to be resolved.** They usually differ because the analyses defined the disorder grouping differently, applied different adequacy rules, or used different time windows. The right response is not to pick one and dismiss the other; it is to state which grouping each describes. A candidate who says "the survey's own figure for common mental disorders, against a later re-analysis restricted to depressive and named anxiety disorders" has demonstrated exactly the discipline the examiner is testing.

11.3 What is deliberately not printed here

The same survey also reports the treatment gap disorder by disorder. Those per-disorder percentages are not printed anywhere in this chapter. They are neither missing nor doubtful: they

are grounded, verified against the survey's own report, and taught, with their derivation and their epidemiological meaning, in the Research Methods, Biostatistics and Epidemiology notes. A per-disorder gap figure is a statement about the epidemiology of that disorder, which is that chapter's subject; an aggregate gap is a statement about access failure, which is this one's.

That boundary is worth stating explicitly rather than silently observing, for a practical reason. If a per-disorder figure feels familiar as you read this, that familiarity is memory rather than evidence, and memory is exactly what the rule exists against. Write the aggregate figures, say that the survey reports the breakdown, and send the reader to the chapter that carries it.

11.4 Demand-side determinants, and the first person a family actually consults

A gap that large is not produced by one barrier; it is produced by a sequence, and the sequence begins before anybody has thought of a clinic. The national survey itself identifies the demand-side conditions behind it: **limited awareness, sociocultural beliefs, values and stigma** are named as demand-side barriers contributing to the treatment gap.

Qualitative work in Indian states fills that in from the point of view of the people making the decision. In a situational analysis of the mental health pathway in one southern state, stakeholders including policymakers, health professionals, non-governmental organisation staff and clients identified stigma, knowledge and expectations, lack of practical treatment access, and cost as influences on decisions about seeking care, drawn from 10 semi-structured interviews and 5 focus groups, each focus group containing 6 to 8 participants.

Stigma is not merely one item on that list; in some populations it is the decisive one. Among 100 consecutive included patients in a Kerala study of people who had attended a faith healer before reaching tertiary psychiatric care, **44 per cent** reported stigma of psychiatric illness as a main reason for avoiding psychiatric treatment.

The pathway data show where people go instead. In a South Indian tertiary outpatient sample of 104 patients and informants, the first carer consulted was an allopathic practitioner for 51.0 per cent, a religious healer for 26.9 per cent, and a mental-health professional for 22.1 per cent. Read those three proportions as a service design brief rather than as a curiosity. Roughly half of first contacts were with a doctor who was not a psychiatrist, which is the strongest available argument for integrating mental health into general and primary care. Around a quarter were with a religious healer, which is the argument for the partnership work taught in this chapter's section on traditional and faith-based care. Fewer than a quarter reached a mental health professional first.

IN INDIA

Design the first contact around where people actually go, not around where you wish they went. In the pathway sample above, 51.0 per cent of first carers were allopathic practitioners and 26.9 per cent were religious healers, against 22.1 per cent mental-health professionals. The service decision that follows is to invest in the recognition and referral capacity of general practitioners and in structured, respectful working relationships with faith-based providers, because both of those sit upstream of any psychiatric service and will continue to whatever the psychiatric service does. Note the limit as you use it: this is one tertiary outpatient sample of 104 patients and informants, not a national estimate.

11.5 The supply side, and the benchmark that does not exist

Demand-side barriers do not act alone. Situational work in northern India identified a dearth of mental-health professionals and inadequate services and facilities as adversely affecting the mental health treatment gap, drawn from 13 purposively recruited interview participants, plus local health-facility records and policy documents.

The workforce position can be put more precisely. For India, **0.29 psychiatrists per 100,000 total population** is reported in WHO Mental Health Atlas 2017, for the reference year 2016, against a denominator of 100,000 of India's total population using the official population estimate shown in the country profile. The definition matters as much as the value: the figure counts a medical doctor who has had at least two years of postgraduate training in psychiatry, which is the Atlas's own workforce definition and not necessarily the definition another source would use.

Now the part that most teaching gets wrong. It is conventional to follow a density figure with a recommended minimum, usually attributed to an international body. No such recommendation could be established for this chapter. An explicit recommended, desirable or minimum number of psychiatrists per lakh of population attributable to the World Health Organization or to the refugee agency was sought across the Atlas report and country profiles for several editions, the global health observatory indicator metadata, the intervention guide, and both agencies' own pages, and it was not found.

The consequence is a refusal that a candidate can make confidently: this chapter prints no workforce benchmark. Attributing a minimum ratio to an international body that has not issued one is a false institutional attribution, and it is worse than printing no benchmark at all, because it borrows an authority that was never given. The defensible sentence is that the density is low in absolute terms, that it is reported by a named body for a stated reference year against a stated denominator, and that no internationally issued minimum could be verified to compare it against.

85.0%	73.6%	0.29	None found
TREATMENT GAP, COMMON MENTAL DISORDERS, OF ACTIVE CASES	TREATMENT GAP, SEVERE MENTAL DISORDERS, OF ACTIVE CASES	PSYCHIATRISTS PER LAKH OF TOTAL POPULATION, ATLAS REFERENCE YEAR 2016	VERIFIED INTERNATIONAL MINIMUM WORKFORCE RATIO

11.6 Distance against money: what actually predicts reaching care

The intuitive model of access is geographical, and the best Indian evidence available here does not support it. In a rural implementation area in Madhya Pradesh, road distance to the nearest public provider of depression treatment was not associated with the probability of seeking treatment: the odds ratio was 1.00, with a 95 per cent confidence interval of 0.98 to 1.02, among 568 adults with probable depression.

In a population-based survey in the same district, what did predict contact coverage was economic position. Among 576 adults with probable moderate to severe depression within a 3,220-adult baseline survey, higher-income occupation was associated with a prevalence ratio of 7.17, with a 95 per cent confidence interval of 1.23 to 41.70, and land ownership with a prevalence ratio of 1.83, with a 95 per cent confidence interval of 1.21 to 2.76 for higher contact coverage.

Read those two findings together and the model of access changes. Building a clinic nearer to people does not, by itself, appear to be the binding constraint; being poorer and owning nothing appears to be. Treat the wide interval on the occupation ratio with the caution it deserves, and note that both studies concern one district and depression rather than the country and mental illness in general. Even with those limits, the direction is a corrective to the standard assumption, and it points service design at cost, lost wages and social position rather than at distance alone.

What is not available is the national quantification of that inequity. Quantified Indian treatment-gap or contact-coverage inequities comparing rural with urban populations, or comparing groups by sex or by education, without resorting to the disorder-specific figures owned by another chapter, were sought in the national survey report, in a summary of it, and in each of the primary Indian studies named above. None was obtained. So the inequity is taught here as a mechanism supported by district-level studies, and no national inequity percentage is printed.

11.7 What closes the gap, and what the trials actually measured

India has produced an unusually strong body of trial evidence on delivering psychological treatment through non-specialists. It is worth knowing well, and it is worth knowing precisely, because almost every trial measured a symptom outcome rather than a coverage outcome, and the distinction is the difference between showing that a treatment works and showing that a gap has closed.

STUDY AND SETTING	RESULT	WHAT WAS MEASURED
Lay health worker led collaborative care in primary care, Goa	A 30 per cent reduction between arms in the prevalence of an ICD-10 common mental disorder diagnosis	A symptom-diagnosis outcome across follow-up, not contact coverage
Brief psychological treatment for depression added to enhanced usual care, Goa	Remission 63 per cent against 48 per cent	Participants meeting the trial's remission criterion; a symptom outcome, not treatment contact or population coverage

STUDY AND SETTING	RESULT	WHAT WAS MEASURED
Brief counselling for harmful drinking added to enhanced usual care, Goa	Remission 54.3 per cent against 31.9 per cent, adjusted prevalence ratio 1.71	A trial-defined remission outcome, not contact coverage
Community champion led psychosocial intervention, Gujarat	Recovery odds ratio 2.2, with a 95 per cent confidence interval of 1.2 to 4.6	Categorical recovery from depression or anxiety symptoms; not service contact or population coverage
Peer-delivered perinatal depression intervention, Goa	Six-month remission 73 per cent against 60 per cent	A trial remission definition assessed six months after childbirth; not contact coverage
Anti-stigma and mental health service delivery in rural primary care, Andhra Pradesh	94.5 per cent of the intervention group against 2.5 per cent of controls self-reported receiving mental-health treatment at 12 months	Self-reported receipt of treatment, not independently verified treatment, symptom remission or whole-population coverage
Grass-roots community programme, rural Maharashtra	Associated with a substantial increase in equitable contact coverage for depression and improved mental health literacy	Contact coverage and literacy indicators before and after implementation

Work down the third column, because that column is the point. Five of the seven measured whether people got better, in populations already identified and enrolled. Only the last two measured whether people were reached at all, and one of those is described qualitatively here because its exact baseline and follow-up coverage figures could not be retrieved from any accessible copy of the original report. So the Indian evidence that non-specialist delivered treatment works is strong. The Indian evidence that any of it has closed a population gap is much thinner, and rests largely on those two.

The control arm figure in the rural Andhra Pradesh trial deserves separate attention. 2.5 per cent of controls self-reported receiving mental health treatment over twelve months, in a high-risk cohort identified by screening. That is the treatment gap seen from inside a trial, in people who had already been identified as needing care, and it is the most concrete illustration in this section of how little happens by default.

TRAP **Citing a remission rate as evidence that an intervention closes the treatment gap.** Remission is measured in people who were enrolled, and enrolment is the very step the gap consists of failing at. An intervention can produce excellent remission in everyone it reaches and change population coverage not at all. When you are asked what closes the gap, name the outcome the study measured before you name the effect size.

11.8 The evaluation nobody has published, and how to answer this

The standard answer to the treatment gap in India is the district programme. An official Indian evaluation reporting a measured before-and-after or controlled change in treatment gap attributable to that programme was sought for this chapter from the ministry, the directorate, the national health mission and the official assessment of the mental health system. None was

obtained: programme objectives, approved districts, administrative activity and service reporting are all documented, and an attributable effect on the gap is not.

That absence is the most useful single thing in this section. The country's principal instrument for closing the gap has not been shown, in any citable official evaluation, to have changed the gap. That is not a claim that it has failed; it is a claim that the question has not been answered, and the reason is that programme reporting was designed to record activity rather than to support attribution.

MNEMONIC

Know, seek, reach, receive, continue

The five steps between having a mental illness and being adequately treated. The gap is the sum of everyone lost at each step, which is why an intervention aimed at one step alone rarely moves the total.

PITFALL

Assuming the gap will fall as services expand. Expansion adds capacity at the reach step and does nothing for the know, seek or continue steps, which is why services can grow while the gap stays where it is. The corollary is practical: if you are given resources to reduce a gap, spend some of them on identification and on retention in care, not all of them on capacity, because capacity that nobody enters is invisible in every outcome measure that matters.

The treatment gap is a proportion of people with the condition, not of the population: 85.0 per cent for common mental disorders and 73.6 per cent for severe mental disorders in India's national survey, and the higher figure belongs to the milder group.

12 · Global and Indian burden of mental disorders: DALYs and YLDs

Counting people with an illness tells a finance ministry almost nothing, because it makes a fatal cancer and a mild phobia the same size. Health accounting solves that by measuring what is lost rather than who is affected, and the instrument it uses is the disability-adjusted life year. This is the measure through which psychiatry competes for resources against cardiology, obstetrics and infectious disease, and a resident who cannot explain how it is built cannot make the case for a mental health service in the room where the money is allocated.

This section teaches the metric and the Indian result. The definitions of incidence, prevalence and the measures of association, and the study designs behind them, are taught in the Research Methods, Biostatistics and Epidemiology notes. The national survey's prevalence findings and the treatment gap have their own sections in this chapter and no figure from either is reproduced here.

12.1 From counting cases to counting lost health

A case count answers how many. It cannot answer how bad, how long, or how much of a life is affected, and those are the questions any allocation decision turns on. Burden metrics were built to make different diseases commensurable by converting everything into a single currency: time lived in less than full health, plus time not lived at all.

The disability-adjusted life year is that currency. **One disability-adjusted life year corresponds to one year of full health lost, and the measure combines years of life lost through premature mortality with years lived with disability**. It therefore adds a severity dimension to a prevalence count, which is exactly what a prevalence count lacks.

Its two components answer different questions and they behave differently.

- **Years of life lost** capture mortality: how many years of expected life a death removed. A disease that kills young generates a great deal of this.
- **Years lived with disability** capture morbidity: how long people live with the condition, weighted by how disabling it is. **One year lived with disability represents one full year of healthy life lost due to disability or ill health, with prevalent conditions weighted by their disability weights**.

RULE **A burden estimate is the product of how many, how disabling, and for how long.** Change any one factor and the total changes without the number of cases moving at all. This is why a condition can become a larger share of national burden while becoming no more common.

12.2 Why psychiatry lives almost entirely in one half of the measure

Mental disorders are, in health accounting terms, a morbidity problem rather than a mortality problem, and that single fact explains most of psychiatry's position in these tables. They begin early, often in adolescence or early adult life. They persist or recur for decades. They impair function substantially while people remain alive. So they accumulate years lived with disability at a high rate and contribute comparatively little through years of life lost.

Three consequences follow, and they are the examinable content of this subtopic.

- **Psychiatry looks small in mortality-based accounts and large in disability-based ones.** Any argument built on death rates alone will understate mental illness, which is why mortality-led health planning has historically neglected it.
- **Duration does the work.** A condition that begins at twenty and persists accumulates burden for decades. This is the quantitative form of the argument for early intervention: preventing or shortening a long illness saves far more health than treating a short one.
- **Disability weights are a modelling choice, not a measurement.** How disabling a condition is judged to be is derived through survey and modelling methods, and a change in weights changes every downstream total without anybody's health having changed. Treat a burden figure as a model output, not as a count.

12.3 The Indian result, and the classification that sits underneath it

The Indian state-level burden work reports the contribution of mental disorders to total health loss in the country and how it has moved. **Mental disorders accounted for 4.7 per cent of total disability-adjusted life years in India in 2017, compared with 2.5 per cent in 1990.** The published estimates carry uncertainty: 2.5 per cent with a 95 per cent uncertainty interval of 2.0 to 3.1 in 1990, and 4.7 per cent with an interval of 3.7 to 5.6 in 2017. The denominator in each case is total disability-adjusted life years in India in the corresponding year, and the estimates were produced for all states of India as part of the global burden of disease work for that cycle.

Now the qualification that changes how the figure should be read, and which most people who quote it have never seen. What is being counted is disability-adjusted life years attributed to mental disorders as defined in that burden cycle, in which suicide was classified under injuries. Self-harm and suicide therefore sit outside this share. Whatever burden India's suicide mortality represents, it is not inside the figure above.

That is not a criticism of the estimate; it is how the classification works, and the same convention is applied across countries so that the comparisons hold. It is, however, decisive for how you use the number. A mental health service argument built on this share is built on non-fatal psychiatric morbidity alone, and a candidate who states that limitation while quoting the figure has demonstrated the four-field discipline more convincingly than one who simply recites the value.

2.5%	4.7%	Not held
MENTAL DISORDERS AS A SHARE OF TOTAL INDIAN DALYS, 1990	THE SAME SHARE IN 2017, SUICIDE CLASSIFIED SEPARATELY	GLOBAL ALL-AGE ESTIMATE IN BOTH DALYS AND YLDS

IN INDIA

When you are making the case for a mental health service to an administrator, use the share-of-burden framing rather than a case count, because it is the currency the allocation decision is already made in: 4.7 per cent of total Indian disability-adjusted life years in 2017, against 2.5 per cent in 1990. Then say, in the same breath, that suicide was classified under injuries in that cycle and sits outside the figure. Volunteering the limitation is what makes the rest of your numbers credible in a room where somebody else is arguing for cardiology.

12.4 How to read a rising share without over-reading it

A share that nearly doubles looks like an epidemic and usually is not one. The share of total burden attributable to a condition can rise for at least five reasons, and only one of them is that more people are becoming ill.

- **Other causes fell.** A share is a ratio. When communicable disease, maternal and child mortality and nutritional deficiency burden decline, everything else rises as a proportion without changing in absolute terms. This is the dominant mechanism in most middle-income countries.

- **The population aged and grew.** More adults living longer means more person-years at risk of conditions that persist.
- **Detection and reporting improved.** Better data produce higher estimates for conditions that were previously invisible.
- **Methods changed.** Successive burden cycles revise disability weights, case definitions and models, and a revision can move a share on its own.
- **Incidence genuinely rose.** Possible, and the hardest of the five to demonstrate.

The disciplined sentence is therefore that the share of India's health loss attributable to mental disorders roughly doubled over that period, that this is consistent with an epidemiological transition in which other causes fell faster, and that a share cannot by itself establish that the underlying rate of illness rose.

TRAP **Reading a rising share of burden as a rising incidence of illness.** A proportion moves when either its numerator or its denominator moves. In a country where the denominator has been falling for other reasons, the numerator can be flat and the share can still climb. Say "share of total burden" every time, and the error becomes impossible to make.

12.5 Three ways to express the same burden

One population, one set of illnesses, three ways to describe them. Each is right for a different audience and each hides something.

EXPRESSION	WHAT IT ANSWERS	WHAT IT HIDES	WHO USES IT
Case count or prevalence	How many people are affected	Severity, duration and mortality; a mild and a severe case count the same	Service planners sizing establishment and workload
Years lived with disability	How much non-fatal health is being lost	Deaths entirely; a condition that kills quickly scores low	Advocates for chronic and disabling conditions
Disability-adjusted life years	Total health lost, fatal and non-fatal together	Who exactly is affected, and the modelling assumptions behind the weights	Finance and planning bodies comparing across disease areas

A complete argument uses all three in sequence: how many people, how much health is being lost, and therefore what share of the allocation is defensible. An argument that uses only the first is easy to dismiss on the ground that the conditions are not serious, and an argument that uses only the third is easy to dismiss on the ground that it is a modelling artefact.

12.6 The global figure this chapter does not print

This item is titled for the global and Indian burden, and no global figure appears in it. That is a declared evidence position rather than an oversight. A global, all-age, both-sex estimate for all mental disorders expressed in both disability-adjusted life years and years lived with disability, with a stated reference year and the corresponding global totals as denominators, was sought

from the international health agency and from the institute that produces the burden estimates, through their observatories and through targeted searches of both. It was not obtained in a form carrying all four fields.

Two further gaps sit alongside it and both concern access rather than existence. A serving address for the published Indian state-level burden analysis could not be reached: the hosted copy returned an error, a repair attempt returned the same dead address, and neither a keyword search nor an exact-title search of the biomedical databases resolved the article. Two separate records in this chapter's evidence ledger point at that same unreachable source. The Indian figures above are printed because they are carried with their full provenance in this chapter's own evidence, not because the article could be opened afresh.

The consequence for a candidate is a specific and defensible position: quote the Indian share with its reference years, its uncertainty intervals, its denominator and its suicide classification; and, if asked for the global figure, say that a global estimate carrying both metrics with their denominators was sought and not obtained in a citable form, rather than reaching for a percentage that is widely repeated and rarely sourced.

MNEMONIC

How many, how bad, how long, how fatal

The four inputs to any burden estimate. Prevalence supplies the first, disability weights the second, duration the third, and mortality the fourth. Psychiatry scores heavily on the first three and lightly on the fourth, which is precisely why it needs this metric to be visible at all.

PITFALL

Mixing burden figures from different estimation cycles in one argument. Successive cycles revise weights, definitions and models, so a figure from one release and a figure from another are not two observations of the same quantity, and the difference between them can be entirely methodological. Take every figure in a single argument from one declared release, and say which one.

Mental disorders accounted for 4.7 per cent of total Indian disability-adjusted life years in 2017 against 2.5 per cent in 1990, a share of total health loss in which suicide was classified under injuries and therefore does not appear.

13 · Basic epidemiological concepts in psychiatry (incidence, prevalence, study designs)

This chapter uses epidemiological measures on almost every page and defines none of them. Incidence, prevalence, the relation between them, relative risk, the odds ratio, attributable risk and the observational designs that produce all of these are taught in the Research Methods, Biostatistics and Epidemiology notes. That is where the definitions, the derivations and the worked arithmetic live, and they are not repeated here.

The division is worth stating precisely because it is unusually clean. Defining a measure is a methods question. Using a measure to decide what a health system should do is a public health question. This chapter is entirely the second, and the three sections immediately before this one are where the using happens: the national survey and its prevalence findings, the treatment gap and its determinants, and the burden of mental disorders expressed as disability-adjusted life years and years lived with disability.

13.1 Where each measure is actually applied in this chapter

A reader who wants to see these concepts working rather than defined should read the three sections named above in that order, because they form a chain and each link uses a different measure for a different decision.

- **Prevalence answers how many people have the condition now.** That is the number a district uses to forecast caseload and to size an establishment. It is applied in the section on the national mental health survey.
- **The treatment gap answers how many of those people are not receiving adequate care.** That is a proportion of the diseased population, not of the general population, and it measures the failure of access rather than the frequency of illness. It is applied in the section on the Indian treatment gap.
- **Burden answers how much health is being lost, weighted by severity and duration.** That is the number a finance ministry compares across diseases when it allocates. It is applied in the section on disability-adjusted life years and years lived with disability.

Read as a chain, those three answer three different questions that are constantly confused with one another in examinations: how common, how untreated, and how costly. They can move in opposite directions. A condition can be common and cheap, or rare and devastating, or common and almost entirely untreated, and each of those combinations implies a different service response.

RULE **Every measure in this chapter is quoted with what it is a proportion of, and that denominator is part of the number.** A prevalence is a proportion of a population. A treatment gap is a proportion of the people with the condition. A burden share is a proportion of total health loss. Three proportions, three denominators, and the commonest examination error in this whole area is carrying one number across to the other's denominator.

13.2 The boundary, and why holding it protects the reader

There is a specific reason this section refuses to give a short definition of prevalence and move on. The measures are written in the same words on both sides of the boundary, so a short restatement here would look identical to correct teaching and would be impossible to police mechanically. What makes it a defect is not the wording but the consequence: two definitions of one measure in one book eventually diverge, and a reader who has met both cannot tell which is authoritative. One owner, one definition.

TRAP **Treating a treatment gap as though it were a prevalence.** They are reported in the same units and in adjacent sentences of the same survey report, and their denominators are entirely different. A gap of the order of four fifths does not mean four fifths of the population is ill; it means that of the people who are ill, roughly four in five are not adequately treated. Say the denominator out loud before you quote the figure.

How common	How untreated	How costly
PREVALENCE, APPLIED TO CASELOAD FORECASTING	TREATMENT GAP, APPLIED TO ACCESS	BURDEN, APPLIED TO ALLOCATION

MNEMONIC

Common, untreated, costly

Three questions, three measures, three denominators, three different service decisions. The measures themselves are defined in the Research Methods, Biostatistics and Epidemiology notes; this chapter only ever uses them.

Incidence, prevalence, measures of association and the observational designs that yield them are taught in the Research Methods, Biostatistics and Epidemiology notes; this chapter applies them and defines none of them.

Suicide prevention and farmer's suicide

14 · Farmer's suicide in India: epidemiology, determinants and prevention

Farmer suicide is the one public health statistic in India that everybody has heard and almost nobody can state correctly. It is quoted in politics, in journalism and in examinations, usually as a bare count and usually without the two things that would make it interpretable: what the counting category actually contains, and what the figure is a proportion of. This section teaches the number properly, and it teaches why the most important quantity in this whole area, the rate at which people in farming die by suicide compared with everybody else, cannot currently be computed at all.

The farming sector as a recognised at-risk group within general suicide epidemiology, and the population prevention levers that apply to it, are taught in the Somatic-symptom, sexual, impulse-control disorders and suicide notes. The national suicide prevention strategy, means restriction as a programme, and the epidemiology of suicide in India generally have their own sections in this chapter. What belongs here is the occupational count itself, its classification, its denominator problem and what would count as evidence that anything has worked.

14.1 Why an occupational category is a public health object

Suicide surveillance in India runs through a police-recorded route. The National Crime Records Bureau (NCRB) compiles cases reported through the police system and publishes them classified in several ways, one of which is by profession. That classification is what turns a distributed

personal tragedy into a countable population, and it is the reason a farming category exists in the national data at all.

An occupational category is worth having for three reasons. It identifies a population that shares an exposure, which is what makes targeted prevention possible. It allows a trend to be watched over time. And it creates political visibility, which in this case has been considerable. It also carries three costs, and the whole of the rest of this section is about them: the category has to be operationally defined, the people in it have to be correctly assigned to it, and the count has to be divided by the right denominator before it means anything.

RULE **An occupational count without an occupational denominator is a workload figure, not a risk figure.** It tells you how many deaths fell in this category. It cannot tell you whether people in that occupation are at higher risk than anyone else, because that comparison requires dividing by the number of people in the occupation, which is a different data source entirely.

14.2 The number, in full

The current published figure is this. **In 2024, 10,546 persons involved in the farming sector died by suicide in India, consisting of 4,633 farmers or cultivators and 5,913 agricultural labourers, accounting for 6.2 per cent of all suicide victims**. The denominator of that percentage is 1,70,746 total suicide victims recorded in India in 2024. The source is the National Crime Records Bureau's annual compilation of accidental deaths and suicides, the reference year is 2024, and the population is persons recorded by NCRB as involved in the farming sector, comprising the farmer or cultivator and agricultural labourer categories.

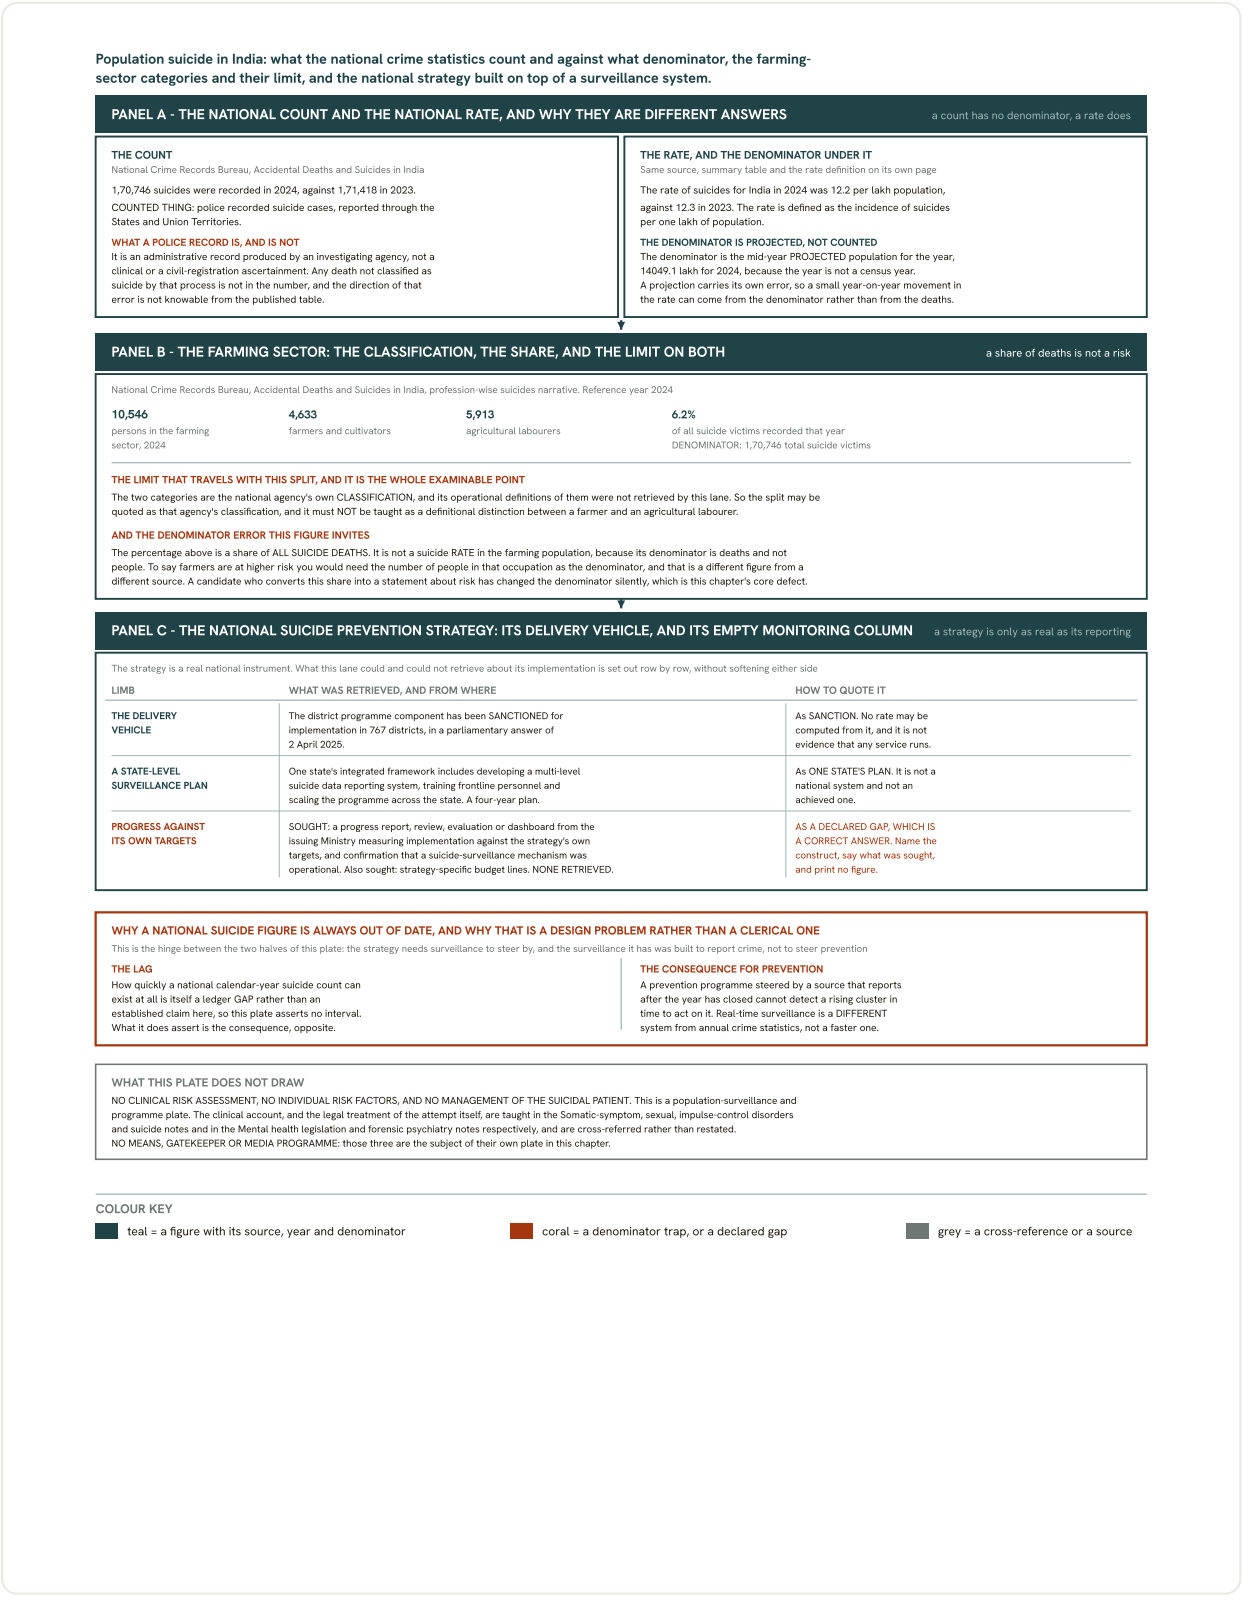


Fig 39.5 Population suicide in India: the count, the rate, the farming-sector categories and the strategy above them. The national agency recorded 1,70,746 suicides in 2024, against 1,71,418 in 2023, at a rate of 12.2 per lakh against 12.3, computed on a projected mid-year population because the year is not a census year. In the farming sector 10,546 deaths were recorded, comprising 4,633 farmers and cultivators and 5,913 agricultural labourers, 6.2 per cent of all suicide victims that year; because the agency's own operational definitions of those two categories were not retrieved, the split is quotable as its classification and not as a definitional distinction, and the percentage is a share of deaths rather than a rate in the farming population. The strategy's delivery vehicle is sanctioned in 767 districts, while progress against its own targets, confirmation that a surveillance mechanism was operational and its specific budget lines were all sought and none retrieved.

The accompanying figure sets out the surveillance architecture this figure comes out of, the categories the bureau publishes and the denominators available for each, alongside the governance of the national prevention strategy. Read it with the two limitations of the next two subtopics in mind, because the figure shows what the system counts and this text explains what the counting cannot support.

10,546	4,633	5,913	6.2%
FARMING-SECTOR SUICIDES RECORDED BY NCRB IN 2024	CLASSIFIED AS FARMERS OR CULTIVATORS	CLASSIFIED AS AGRICULTURAL LABOURERS	OF ALL RECORDED SUICIDE VICTIMS, NOT OF PEOPLE IN FARMING

Two features of that set deserve immediate attention, because both are routinely lost. The larger of the two component groups is agricultural labourers rather than landowning cultivators, which cuts against the way the issue is usually framed in public discussion. And the percentage is a share of all suicides in the country, which is a statement about the composition of the national total and not about risk in the farming population.

14.3 The classification, and the definition that could not be obtained

The split between the two categories is the most quoted feature of this data and the one this chapter is least able to explain. The bureau publishes the farming sector total divided into farmers or cultivators and agricultural labourers, and this chapter prints that split because it is what the source prints. What could not be obtained is the bureau's own operational definition of either category: the rule a recording officer applies to decide which of the two a particular deceased person belongs in.

That absence sets a hard limit on what may be said, and the limit is worth stating as a rule of method. The split may be reported as the bureau's classification. It may not be taught as a definitional distinction between two groups, because exactly what separates them in the bureau's coding is the thing that was not obtained. Anyone who tells you that the categories divide people by landholding, by tenancy status or by whether the work is paid is supplying an operational definition from inference, and the inference may be right, but it is not sourced.

This matters practically and not just pedantically. If a service or a state government wants to target the larger group, it needs to know who is in it. A category whose boundary rule is unpublished cannot be used to design an intervention with any confidence about whom it will reach.

14.4 The denominator problem, which is the heart of this topic

Here is the central analytical point of the section, and the one most worth carrying into an examination. To say that people in farming are at elevated risk of suicide, you need a rate: deaths in the farming population divided by the size of the farming population, compared with the corresponding rate elsewhere. The numerator exists and is published. The **denominator** does not exist in any form that can be validly matched to it.

The reason is that the two quantities are produced by different systems that do not share a definition of who is a farmer.

- **The numerator is an administrative classification made at the time of death,** by a recording officer, using a coding rule that this chapter could not obtain, applied to whatever information about the deceased was available at that moment.
- **Any candidate denominator is a census or survey classification,** made under its own definition of agricultural work, in a different reference period, and with its own treatment of people who do several kinds of work.
- **Neither system was designed to be divided by the other.** Combining them produces a rate whose numerator and denominator describe overlapping but non-identical populations, which is the classic route to a fabricated statistic that looks entirely respectable.

So the published share, being of all suicide victims rather than of people in farming, is doing something quite different from what most readers assume. It tells you what proportion of India's recorded suicides occurred among people classified as being in farming. It cannot tell you whether being in farming raises the risk, and it will move if the total number of suicides in the country changes even when nothing whatever happens in agriculture.

TRAP **Converting the share of all suicides into a statement about risk in farmers.** A share of the national total has the national total as its denominator, so it responds to changes anywhere in the country. Saying that farming accounts for a given percentage of suicides is correct. Saying that farmers are therefore at that level of risk, or comparing the share with the proportion of the workforce in agriculture as though the comparison were a risk ratio, is an error of denominators dressed up as an analysis.

IN INDIA

If you are asked to design a district response, do not begin from the national share. 6.2 per cent of all recorded suicide victims is a composition figure whose denominator is 1,70,746 total suicide victims recorded in India in 2024, and it tells you nothing about your district's farming population. Begin instead by establishing what your local denominator would be and who defines it, because until that is settled you cannot tell whether an intervention has worked, only whether the count went up or down for reasons you cannot attribute.

14.5 Who gets counted as a farmer, and who does not

Even before the denominator problem, the numerator has an assignment problem. Occupational classification at the time of death rests on what somebody records about a person's work, and agricultural work in India is unusually hard to record cleanly.

PERSON	HOW THE DEATH MAY BE CLASSIFIED	WHY THE ASSIGNMENT IS UNCERTAIN
An owner-cultivator working land held in their own name	Most likely within the farming sector	The clearest case, and the one public discussion imagines as typical
A tenant or sharecropper working land held by somebody else	Could fall in either component, or outside the sector	Whether cultivation without ownership counts as cultivating depends on a coding rule that is not published
A woman working on family land without holding it	Frequently recorded by household role rather than by occupation	Unpaid family labour is systematically under-recorded as work in many administrative systems
A person who farms in one season and does other paid work in another	Depends on which occupation the record captures	Multiple occupation is common and the record usually accommodates one
A landless labourer working across several holdings	Most likely the agricultural labourer component	Casual and migratory work is the hardest to attribute to any single occupation

Read that table as a description of measurement error with a direction rather than as a list of exceptions. The cases that are hardest to classify are concentrated among the poorest, the landless, the migratory and among women. Measurement error that is heavier at one end of a social gradient does not simply add noise; it biases the picture, and it biases it towards the group that is easiest to identify, which is the landholding cultivator.

14.6 What the data can and cannot say about causes

Determinant claims about this population are made constantly and are rarely qualified. Two limits govern how far a police-recorded surveillance system can support them.

The first is that a cause field in an administrative record captures the reason recorded by an investigator, from information available shortly after a death, in a context where the family may have strong reasons to present some explanations and not others. It is a recorded attribution, not an aetiological finding, and treating it as one imports the recorder's framework into what looks like epidemiology.

The second is that the same denominator problem prevents any comparison of determinant frequency between the farming population and everybody else. Without a rate, a determinant that appears often in this group cannot be shown to appear more often than in another group, which is what a risk factor claim requires.

The determinants of agrarian distress themselves, and the general aetiological models of suicide within which they sit, are taught in the Somatic-symptom, sexual, impulse-control disorders and suicide notes, and are not re-established here. What this section adds is the measurement caution that has to travel with them.

14.7 Prevention, and what would count as evidence

Targeted prevention for this population has been attempted in several forms, and this chapter holds no evaluation of any of them: no controlled comparison, no before-and-after estimate with a stable denominator, and no published attribution of change in the count to any programme. That absence is worth teaching directly, because it lets a candidate say something more useful than a list of interventions.

Set out what an evaluation would need, and the difficulty becomes obvious.

- **A stable denominator** over the evaluation period, so that a falling count can be distinguished from a shrinking population working in agriculture.
- **A comparison group** not exposed to the intervention, since counts in this area move with agricultural, economic and reporting conditions that no programme controls.
- **A consistent classification rule** across the period, because a change in coding practice would move the count on its own.
- **A pre-specified outcome and period,** since with year-to-year variation a favourable window can be found after the fact in almost any series.

Population-level prevention levers that apply here, including restriction of access to means and the national strategy that organises them, are taught in their own sections of this chapter, and this section does not restate them.

MNEMONIC

Who counted, who was counted, out of what

Three questions for any occupational suicide statistic. Which system produced the count, which coding rule assigned people into the category, and which denominator the figure is a proportion of. For this statistic the first is answerable, the second was not obtained, and the third does not exist in matchable form.

PITFALL

Comparing this year's count with last year's and calling the difference a trend. Year-on-year movement in an administratively recorded count reflects changes in reporting and classification practice as well as changes in deaths, and neither is separable without a stable denominator and a stable coding rule. Report the count with its year and its source; say that a trend claim requires both of those to have held constant, and that this chapter cannot establish that they did.

The farming-sector figure is a count with a national denominator, not a rate: 10,546 recorded farming-sector suicides in 2024, being 6.2 per cent of 1,70,746 total recorded suicide victims, and no valid occupational denominator exists to turn it into a risk.

15 · National suicide prevention strategy and population-level suicide prevention approaches

A national strategy is a governance document, not a treatment. It states what a country intends to do about a cause of death, names the sectors that will do it, and sets targets against which somebody is later supposed to measure it. That last clause is where most of the examinable content of this topic lives, because it is the clause that most often goes unhonoured. Candidates are rarely asked to recite a strategy's contents; they are asked what a strategy is for, what has to exist underneath it before it can work, and how one would know whether it had. Those are service questions, and this section answers them.

The prevention framework itself, the tiering of interventions into universal, selective and indicated, the four evidence-based population interventions in the World Health Organization's implementation guidance, the strategy's own headline goal and its staged action horizons, is taught in full in the Somatic-symptom, sexual, impulse-control disorders and suicide notes and is not restated here. What this section adds is the layer above and below that framework: how a strategy is delivered through an existing programme, what a surveillance system has to be able to do before a target can be checked against it, what the financing question is, and, unusually for a textbook section, a candid account of what is not published at all.

15.1 What a population-level strategy actually is, and what it is not

India has a **National Suicide Prevention Strategy**, which is the national governance instrument for suicide prevention and the reference point for every state and district plan written after it. A resident should be able to say three things about any such document. First, it is **multisectoral** by necessity: the levers that move population suicide mortality sit largely outside health, in agriculture, in commerce, in education and in the media, and a strategy that binds only the health ministry has already conceded most of its own case. Second, it is *staged*, because the levers act on different timescales; a regulatory restriction can act within a year, while a schools programme acts over a generation. Third, and this is the part candidates skip, it is *accountable*, meaning it names targets and therefore commits somebody to measuring them.

What a strategy is not is a service. It does not by itself open a clinic, train a worker or restrict access to anything. Everything it promises must be delivered through machinery that already exists or must be built, and in India that machinery is overwhelmingly the district-level mental health programme sitting under the national programme. This is why a question about the national strategy is very often, in substance, a question about district implementation.

RULE **A strategy is judged by its delivery vehicle and its measurement system, never by its text.** Two countries with identical strategy documents will produce different mortality curves if one has a functioning district service and a timely surveillance system and the other has neither. When you are asked to appraise a national strategy, appraise those two things first and the document second.

15.2 The delivery vehicle, and the single most misread figure in Indian community psychiatry

The vehicle for population suicide prevention in India is the **District Mental Health Programme**, the district-level component of the national programme. Its administrative scale is published: the district programme has been sanctioned for implementation in **767** districts. That is a governmental statement of approval, made in reply to a question in Parliament, and it is the number a candidate will meet on a slide.

Read what it counts. It counts districts for which implementation has been *sanctioned*. It does not count districts in which a suicide-prevention service exists, in which a gatekeeper has been trained, or in which a person in crisis could be seen. Sanction is an act of the centre; function is a property of the district. The two are separated by staffing, by the release of funds, by the state's own absorption capacity and by time, and nothing in a sanction figure carries information about any of them. Nor does the figure support a rate: there is no valid denominator against which a count of approved districts becomes a proportion of anything meaningful, because approval is cumulative and administrative rather than a measurement of a population at a point in time.

TRAP **Reading a sanctioned-district total as coverage.** Candidates write that the programme "covers" the sanctioned number of districts and then reason from that to national access. The word in the source is approval, and the correct sentence in an answer is that the programme is sanctioned in that many districts while the number in which suicide-prevention services actually function is not established. That second half of the sentence is worth more marks than the number is, because it shows you know what the number measures.

IN INDIA

When you are asked what a district psychiatrist can actually do under the national strategy, answer through the district programme, because that is the only route the strategy has. The practical consequence is that a district's suicide-prevention capability is decided by whether its programme post is filled and its funds released, not by whether the district appears on the sanctioned list. A resident asked to plan district suicide prevention should begin by establishing which of the sanctioned components are staffed and running locally, and treat the national total as a statement about paperwork rather than about services.

15.3 Surveillance read as an instrument, not as a table of numbers

A strategy that sets a mortality target needs a measurement system capable of telling it whether the target moved. In India **suicide surveillance**, meaning the routine national counting of deaths by suicide for programme purposes, is produced by the police-reporting system and published annually as a compendium of accidental deaths and suicides. Treated as a data instrument rather than as a source of quotable totals, it has three properties that decide what it can and cannot be used for.

- **Case definition.** The unit is a case recorded by police, which is an administrative event with legal consequences and not a clinical determination. Anything that changes the willingness of families or of the police to record a death as suicide changes the series without anything changing in the population.

- **Timing.** The annual publication collates a full calendar year, so collection can only begin once that year has ended. The interval between an event and its appearance in the national series is therefore at least a year and in practice longer.
- **Granularity.** The series is compiled upward from states, so its usefulness for targeting depends on the quality of state-level reporting, which is not uniform.

The timing property is the one with programme consequences and it is worth stating plainly as a declared limitation rather than glossing it. A system whose national picture arrives a year or more after the events it describes cannot function as an early-warning instrument, cannot support the rapid evaluation of a newly introduced intervention, and cannot tell a state within a planning cycle whether what it did last year worked. Whether India now has a suicide-surveillance mechanism that is operational, nationally complete and timely enough to monitor strategy outcomes is not established: this chapter sought an official assessment of exactly that question through the health ministry, the bureau that publishes the national series and the national institute that advises on the programme, and none was retrieved. The honest examination answer is that the strategy names a surveillance objective and that no published assessment confirms the objective has been met.

Sanction	None	Over a year	Not established
WHAT THE DISTRICT TOTAL COUNTS, NOT SERVICE DELIVERY	VALID DENOMINATOR FOR THE DISTRICT TOTAL, SO NO RATE	MINIMUM LAG BEFORE A CALENDAR YEAR ENTERS THE NATIONAL SERIES	PROGRESS MEASURED AGAINST THE STRATEGY'S OWN TARGETS

15.4 The state as the real implementing unit

Health is administered by the states, so a national strategy is realised as a set of state strategies and state programmes, and the interesting variation is between states rather than around the national mean. At least one state has taken this seriously enough to build an integrated framework of its own for suicide prevention and response, spanning health workers, communities and law-enforcement agencies. Its programme is designed around three components that are worth learning as a template: **a multi-level suicide data reporting system, training for frontline personnel, and scale-up of the programme across the whole state**.

Notice the order and the logic. The reporting system comes first because without it the state cannot know its own baseline or detect change; training comes second because it converts the strategy into an act somebody performs; scale-up comes third because a pilot in two districts answers no population question. That sequence is the generalisable teaching point, and it is a better answer to "how would you implement the national strategy in your state" than any recitation of the national document.

The specifics of any individual state's instrument, including the date on which such a strategy was released and the contents of its published text, are not established in this chapter's sources and should not be reconstructed from a seminar handout. Name the state layer, describe the template, and say that the state documents themselves are the primary sources a candidate would need to read.

15.5 Financing, which is where most strategies are actually decided

A strategy without money attached is a statement of intent. This is not a cynical observation; it is the standard first question in any appraisal of a public-health programme, and it has been raised officially in India in relation to both the mental health legislation and the suicide prevention strategy, where the absence of dedicated budgetary allocation has been identified as a concern about implementation. An early official parliamentary assessment recorded that whether the strategy was being implemented effectively remained undetermined.

What this chapter cannot tell you, and says so rather than inventing it, is how much money the strategy has. Allocations, releases and utilisation specific to the suicide prevention strategy, separated from the broader mental health programme and from the national institutions, were sought from the health ministry, from budget documentation and from parliamentary answers, and were not retrieved. Two things follow for an examination answer. First, do not quote a rupee figure for the strategy; there is none you can defend. Second, the fact that a national strategy's own financing cannot be isolated from the wider programme is itself a finding about the strategy, and saying so demonstrates more understanding than a number would.

> **PITFALL**
>
> Treating an allocation as an outcome. Where a programme's spending is published, clinicians frequently quote the allocation as evidence of activity. An **allocation** is a budget line, a **release** is money that has left the treasury, and **utilisation** is money that has been spent; the three routinely differ, and none of them is a service. If you use a financial figure at all, say which of the three it is, and for which year.

15.6 Evaluating a strategy against its own targets

The cleanest way to appraise any national strategy is to hold it to the standard it set for itself, because that standard was chosen by the authority that wrote it and cannot be dismissed as an outsider's benchmark. That approach requires a progress report, a review, an evaluation or a live dashboard that measures implementation against the strategy's stated targets, published by the issuing body or by an independent auditor.

No such assessment was retrieved for the Indian strategy. The search covered the health ministry's own programme pages and its published strategy location, parliamentary answers in both houses, a parliamentary standing committee report, government press releases and the national institute's own domain. The absence is the finding, and it should be taught as one: India has a national strategy, a delivery vehicle with a published sanction footprint, a surveillance system with a structural reporting lag, and no published measurement of the strategy against its own targets. A candidate who can lay out those four statements in order has given a better account of Indian suicide prevention than one who recites the strategy's contents.

It is worth setting the four appraisal questions out against each other, because the pattern that emerges is the section's real teaching point. In every case the design statement is published and traceable to an authority, and in every case the delivery statement is either missing or is a different quantity wearing the design statement's clothes.

APPRAISAL QUESTION	WHAT IS PUBLISHED AS DESIGN	WHAT DELIVERY WOULD REQUIRE	STATE OF THE DELIVERY EVIDENCE
Reach	Districts for which the district programme is sanctioned, stated by the health administration in reply to Parliament	Districts in which a suicide-prevention service is staffed and seeing people	Not established; sanction and function are different quantities and only the first is published
Measurement	A surveillance objective inside the strategy, and an annual national compendium of recorded suicides	A system that is nationally complete and timely enough to detect change within a planning cycle	Not established; the annual series can only begin collecting a calendar year once it has closed
Money	The mental health programme's own budget lines	Allocation, release and utilisation attributable to the strategy itself	Not retrieved; strategy-specific financing cannot be separated from the wider programme
Effect	Targets stated by the strategy for itself	A progress report or evaluation from the issuing body or an auditor, measured against those targets	Not retrieved from the ministry, from Parliament, from a standing committee or from the national institute

Two features of that table are examinable in their own right. The first is that the left-hand column is answerable for almost any country and the right-hand column is answerable for very few, which is why international comparison of suicide-prevention strategies is so much weaker than it looks. The second is that a candidate who has learned only the left-hand column will answer a delivery question with a design statement and will not notice the substitution, because the design statement is a real, sourced, governmental figure. The error is not one of accuracy but of scope, and scope errors are invisible from the inside.

MNEMONIC

Vehicle, Surveillance, Sectors, Money, Measurement

The five questions to ask of any national suicide prevention strategy, in the order in which they decide whether it can work. What delivers it, what measures it, who outside health is bound by it, what funds it, and who checks it. Every one of the five is answerable for India, and two of the answers are that nothing is published.

15.7 Writing about this in an examination

Three habits separate a strong answer from an average one here. The first is scope discipline: every national figure you print must travel with the body that issued it, the period it refers to, the population it describes and what it is a proportion of, and if you cannot supply all four you should describe the thing without its number. The second is the distinction between design and delivery, which recurs in every subtopic above: a sanctioned district, a stated objective and an allocated budget are all statements of design, while a functioning service, a measured outcome

and spent money are delivery. The third is the willingness to declare an absence. In a viva, "that figure is not reliably published, and here is what was looked for" is a stronger answer than a confidently quoted number of unknown provenance, and it is the answer this chapter is built to let you give.

The other suicide-prevention material in this chapter divides cleanly. The specific population interventions, means policy, gatekeeper programmes and media guidance, are configured and evaluated in their own section; the agrarian category argument and its denominator problem belong with farmer suicide; the national data series read as a distribution belongs with suicide epidemiology. The clinical prevention framework, the intervention tiers and the international implementation guidance behind all of them belong to the Somatic-symptom, sexual, impulse-control disorders and suicide notes.

> A national suicide prevention strategy is only as real as its delivery vehicle, its surveillance timeliness and its financing, and in India the first is sanctioned in a published number of districts, the second reports a calendar year only after that year has closed, and the third cannot be separated out at all.

16 · Means restriction, gatekeeper training and media guidelines in suicide prevention

Three interventions, and each of them is two different things. Means restriction, gatekeeper training and media guidance can each be described as something a clinician does with a person in front of them, and each can be described as something a state designs, funds, delivers and evaluates across a population. The clinical readings are taught elsewhere and are cross-referred below. This section takes the second reading throughout, because that is where the marks in a public-health paper sit and because the two are genuinely different objects: a gatekeeper programme is commissioned and evaluated, a gatekeeper referral is made.

The mechanisms behind the three, why restricting access to a method works when crises are short and ambivalent, what the evidence says about gatekeeper training's effect on recognition and on mortality, and the paired imitation and protective effects of media coverage, are taught in full in the Somatic-symptom, sexual, impulse-control disorders and suicide notes. Means restriction as a task carried out with a young person and their family, including the household conversation and the safety plan it belongs to, is taught in the Child behavioural, emotional and other disorders notes. Neither is restated here. What follows is the Indian implementation layer: which instruments exist, who issues them, what has actually been measured, and what has not.

16.1 The three levers, read as programmes

A population lever has four parts that a clinical act does not: an issuing authority, a legal or administrative instrument, a delivery route, and an evaluation endpoint. Run the three interventions through those four parts and the differences between them become obvious.

- **Means restriction** is issued by a regulator, not by a health department, and its instrument is usually a statute or a subordinate order. Its delivery route is the market and the supply chain

rather than the clinic, which is why it reaches people who never present to a service. Its endpoint is method-specific mortality.

- **Gatekeeper training** is issued by a health or education authority, its instrument is a curriculum, its delivery route is a workforce, and its endpoint is contested: recognition and referral are measurable, whereas an effect on deaths is not demonstrable from a training programme alone.
- **Media guidance** is issued by a press or broadcast regulator, its instrument is a code of conduct, its delivery route is editorial practice, and its endpoint is the content of reporting itself rather than any clinical event.

That table of four parts is the frame for everything below, and it is also the answer to the commonest structural mistake in this topic, which is to compare the three on effectiveness as though they were competing treatments. They are not substitutes. They act on different populations through different institutions, and a national programme runs all three because each reaches people the other two miss.

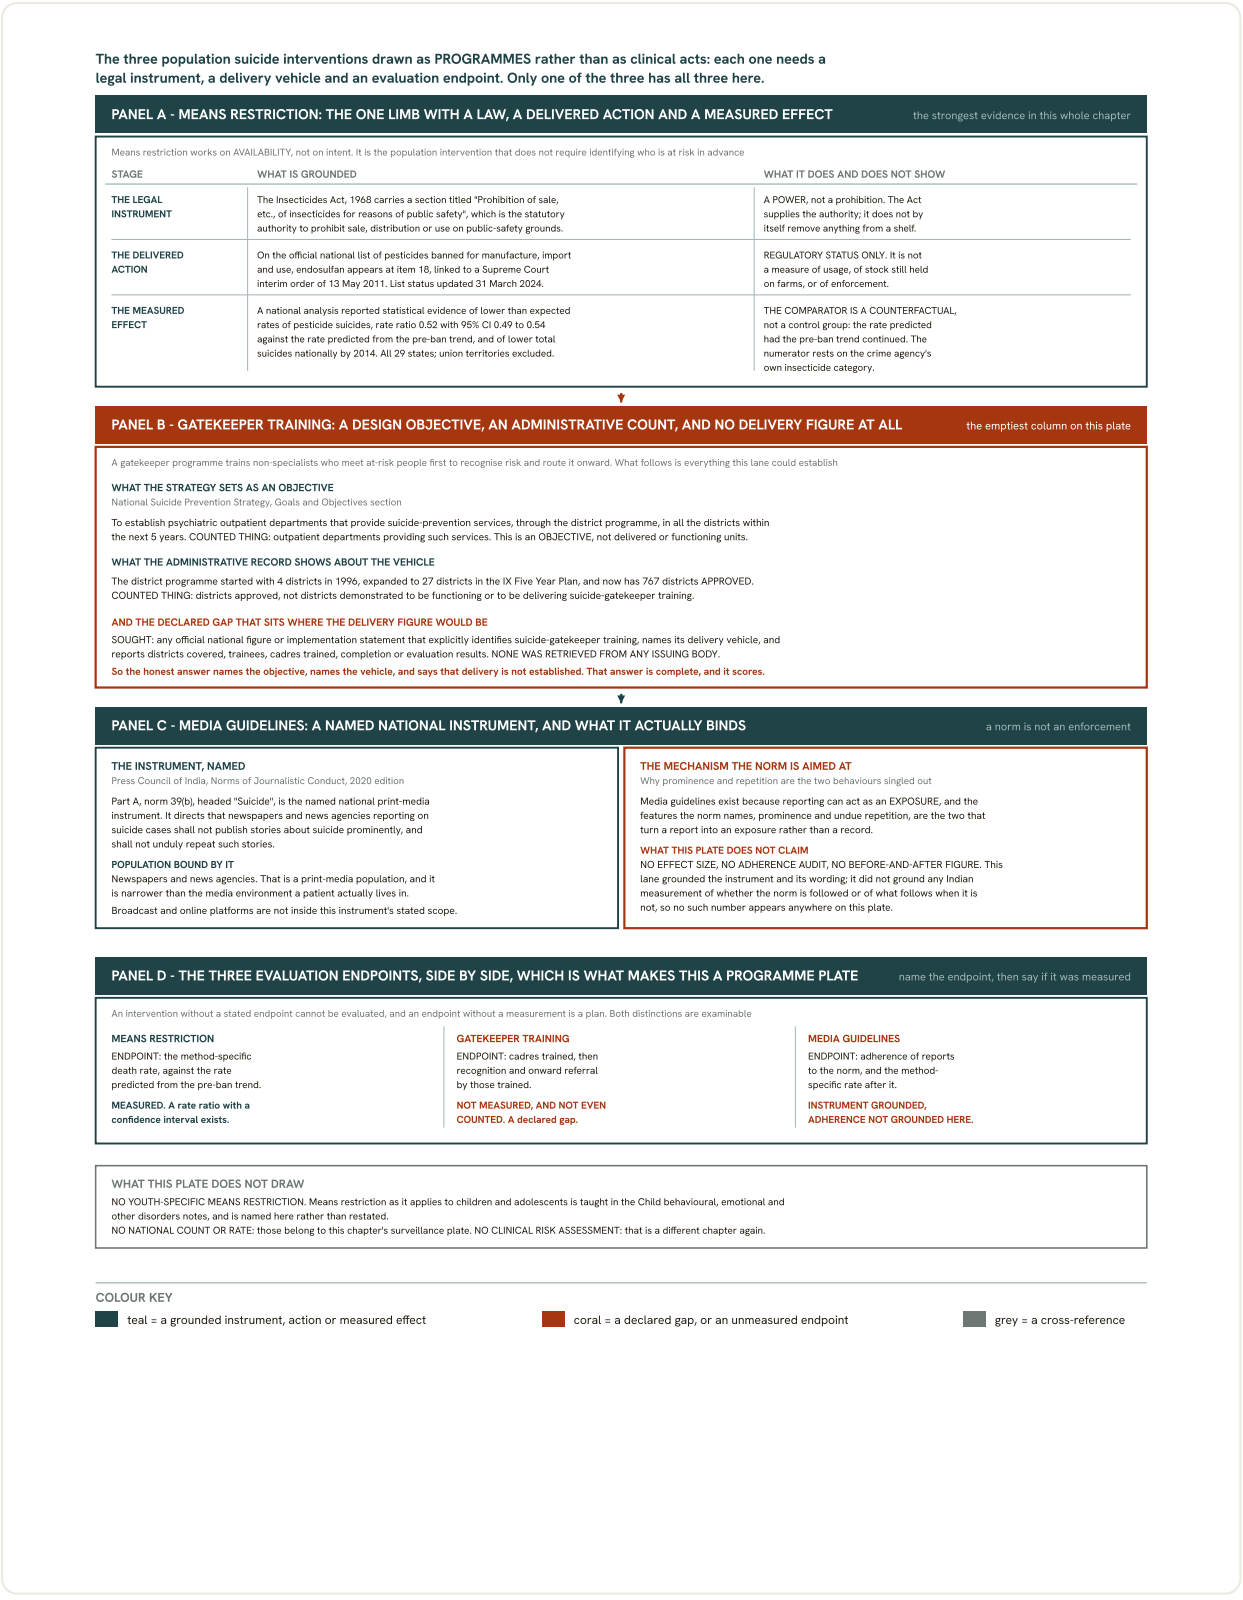


Fig 39.6 Means restriction, gatekeeper training and media guidelines, each drawn as a programme with its own evaluation endpoint. Means restriction is the one limb with all three stages grounded: a statutory power to prohibit an insecticide for reasons of public safety, a delivered ban listing endosulfan at item 18, linked to a Supreme Court interim order of 13 May 2011, and a measured effect, a pesticide-suicide rate ratio of 0.52 with 95% CI 0.49 to 0.54 against the rate predicted from the pre-ban trend. Gatekeeper training has an objective set by the national strategy and a vehicle approved in 767 districts, but no official figure naming trainees, cadres, completion or evaluation was retrieved from any issuing body. The media limb has a named national instrument directing that newspapers shall not publish suicide stories prominently or unduly repeat them, bounded to print media, with no Indian adherence measurement grounded here.

RULE **Name the issuing authority before you discuss the intervention.** A population lever that no institution is empowered to pull is a wish. In India the three levers sit with three different authorities, only one of which is the health ministry, and half of the implementation difficulty in this topic follows from that single fact.

16.2 Means restriction in India: which instruments actually exist

India regulates agricultural chemicals through a dedicated statute. The **Insecticides Act, 1968** carries a provision, at Section 27, headed as a prohibition on the sale and related dealings in an insecticide for reasons of public safety. Read as a public-health instrument rather than as an agricultural one, that is the legal hook on which access restriction hangs: it is the power to take a compound off the market, and it exists independently of any mental health law.

The operational expression of that power is a published national list of pesticides that are banned, refused registration or restricted in use, maintained by the plant-protection directorate. Compounds move onto it by different routes, including judicial ones: **endosulfan appears in the banned-for-manufacture, import and use section of that list at item 18, linked to an interim order of the Supreme Court made in 2011**. That detail is worth carrying because it shows the mechanism by which access restriction is achieved in practice, which is administrative and judicial rather than clinical, and because it locates the document a resident should cite.

It is worth pausing on why the instrument is agricultural rather than medical, because the point generalises. Population access restriction almost always runs through the law that governs the commodity, not through the law that governs health: a pesticide through an insecticides statute, a medicine through drug scheduling and pack regulation, a structure through building and transport rules. The health system's role in each case is to supply the mortality evidence that justifies the restriction and then to measure what follows, not to issue the restriction itself. A resident who understands that will not waste an answer proposing that a mental health authority ban anything, and will instead describe the correct route, which is evidence to the sectoral regulator and, where that fails, to the courts.

Two cautions travel with any statement about that list. It records *regulatory status*, not use and not enforcement; a compound can be banned on paper and remain in circulation in a district. And a list is a statement about compounds, not about deaths, so nothing about mortality follows from the list itself. The mortality question is separate and is answered below.

16.3 What the regulation bought, and how anyone knows

Regulation of a method is one of the few population interventions in psychiatry that can be evaluated at national scale, because the exposure changes on a known date for the whole country and the outcome is recorded routinely. The Indian evaluation used exactly that design. Comparing the observed rate of pesticide suicides against the rate predicted by extrapolating the stable pre-ban trend, a national analysis reported **a pesticide-suicide rate ratio of 0.52, with a 95 percent confidence interval of 0.49 to 0.54**, together with a finding that total suicides nationally were lower than expected by 2014.

Learn the structure of that claim, not just its value. The comparator is a counterfactual, a modelled rate rather than an observed one, which means the estimate inherits every assumption in the trend model. The outcome is drawn from routine national records of deaths classified as pesticide poisoning, and the way those record categories are defined, and what they can and cannot bear, is the argument taken up where farmer suicide is taught in this chapter; it is not repeated here. The second finding, that total suicides also fell below the predicted line, is the analytically interesting one, because it is the finding that speaks to whether restriction merely displaces the method.

PITFALL

Quoting a rate ratio as though it were a proportion. A ratio below unity from a counterfactual analysis says the observed rate was about half of what the pre-ban trend predicted for that year. It does not say half of suicides were prevented, it does not say half of pesticide deaths stopped, and it cannot be converted into a number of lives without the predicted rate and the population it applies to. State it as a ratio against a predicted rate, or do not state it at all.

16.4 Gatekeeper programmes: what has to be specified before one exists

A **gatekeeper** programme is a training intervention aimed at people who meet those at risk before any clinician does. As a programme rather than as a course, it is defined by six specifications, and a plan that omits any of them is not yet a programme: which cadre is trained, what competency the training is supposed to produce, who delivers and refreshes the training, what the trainee is expected to do when they identify somebody, what receives the referral, and what will be counted afterwards. The sixth is the one most often left blank, and it is the one that decides whether the programme can ever be shown to have worked.

India's obvious delivery routes for such a programme are the district mental health programme and the national tele-counselling service, both of which already reach a workforce and a public. What is *not* available is any official national statement that identifies suicide-gatekeeper training as a named component, names its delivery vehicle, and reports districts covered, personnel trained, cadres included, completion rates or evaluation results. That was sought from the health ministry, the health services directorate, the national health mission and the bodies responsible for programme implementation, and it was not retrieved. Say so in an answer, because it is the truthful position and because it explains why the Indian literature on gatekeeper training consists of individual studies rather than programme reports.

Choosing that endpoint honestly is harder than it looks. Recognition and referral behaviour are the endpoints a training programme can plausibly move and can measure within a reasonable period, and they are legitimate. Deaths are the endpoint everybody wants, and a single training programme is almost never powered to detect a change in them, because the event is rare in any one district and because a trained workforce is only one link in a chain that also requires a service to refer into. Writing mortality into the evaluation plan of a district training programme is therefore not ambition; it is a design error that guarantees a null result and discredits a

reasonable intervention. State the intermediate endpoint, state that it is intermediate, and say what would have to be true downstream for it to matter.

TRAP **Treating an approved district as a trained district.** The national administrative scale of the district programme is published, and it is regularly imported into answers about gatekeeper coverage. Approval of a district under the programme carries no information at all about whether anybody in that district has been trained to recognise suicide risk. Approval is an accounting fact about the centre; training is a fact about a workforce, and no published national count of the second exists.

16.5 Media guidance: the instrument, and the reach it does not have

A **media guideline** is a code binding on publishers, and in India the named national instrument for print is the Press Council of India's **Norms of Journalistic Conduct, 2020**, which contains a suicide-specific norm directed against reporting suicide prominently and repeating such stories unduly. Two features of that instrument matter more than its contents.

The first is its scope. It is a code for newspapers and news agencies, which is to say for print journalism and the agencies that feed it. Broadcast media and the online and social platforms on which most people now encounter a suicide story sit under different regulators or under none, so compliance with a print code is not compliance across the media environment. The second is its force. A council code operates by adjudication and censure rather than by penalty, so its effect depends on editorial willingness in a way that a pesticide ban does not.

The practical consequence for a resident is that the evaluation endpoint for media guidance is the content of coverage, measured against the code's own requirements, rather than any mortality outcome. If you are asked how you would evaluate an Indian media-guidance intervention, the answer is a content audit of reporting before and after, sampled across the outlets the code actually binds, and an honest statement that the outlets it does not bind are the ones carrying most of the audience.

16.6 Delivering all three through one district system

Whatever the national strategy says, everything it promises arrives in a district through a single service. The strategy addresses this directly: **the 2022 strategy sets, as a design objective, the establishment of psychiatric outpatient departments providing suicide-prevention services through the district mental health programme in all districts within the next five years**. That is a target, and reading it as an achievement is the single commonest error in this topic.

Against that target sits the administrative position: the health services directorate reports the district programme's national scale as **767** approved districts. Approval is what that number counts. It is not a count of districts with a functioning suicide-prevention outpatient department, and it must not be rewritten as gatekeeper delivery or as any other service output. The gap between a design objective expressed in all districts and an approval figure expressed in a count of districts is not a contradiction; the two are simply different quantities, and an answer that treats them as one has made a scope error rather than an arithmetic one.

Regulator	Press council	District programme	Not established
ISSUING AUTHORITY FOR MEANS RESTRICTION IN INDIA	ISSUING AUTHORITY FOR THE PRINT MEDIA CODE	DELIVERY ROUTE THE STRATEGY NAMES FOR ALL DISTRICTS	NATIONAL GATEKEEPER TRAINING COVERAGE, ANY CADRE

16.7 The evidence state of each lever, side by side

The most useful thing a candidate can carry out of this section is not a set of figures but a map of what is and is not established for each lever in India. The pattern is asymmetric in an instructive way: the lever with the strongest legal instrument is also the only one with a national effect estimate, and the lever that is easiest to run is the one with no national data at all.

LEVER	INDIAN INSTRUMENT, AND WHO ISSUES IT	WHAT IS ESTABLISHED ABOUT EFFECT	WHAT IS DECLARED NOT ESTABLISHED
Access restriction	A dedicated insecticides statute with a public-safety prohibition power, plus a national banned and restricted list maintained by the plant-protection directorate	A national counterfactual analysis reporting a pesticide-suicide rate ratio below unity against the predicted pre-ban trend, with a total-suicide finding in the same direction	Enforcement and actual availability at district level, which regulatory status does not measure
Gatekeeper programmes	No named national instrument retrieved; the district programme and the national tele-counselling service are the available delivery routes	Nothing at national level	Districts covered, personnel trained, cadres included, completion and any evaluation result
Media guidance	A press council code of journalistic conduct with a suicide-specific norm, binding on newspapers and news agencies	Nothing at national level as an outcome; the endpoint is compliance in content	Compliance auditing, and coverage of broadcast and online outlets the code does not bind

IN INDIA

The operational decision this section forces is which lever a district team can actually pull. Access restriction is not theirs: it sits with an agricultural regulator and with the courts, so the district contribution is surveillance and advocacy rather than action. Media guidance is not theirs either, beyond a local relationship with editors. Gatekeeper training is the one lever a district service can commission itself, which is precisely why it is the lever most often run without a specified endpoint. If you are asked to plan district suicide prevention, commission the training you control, and write the count you will report before the first session is delivered.

MNEMONIC

Authority, Instrument, Route, Endpoint

The four parts of any population lever, and the order in which to write them. Missing authority makes the intervention unenforceable; missing instrument makes it voluntary; missing route makes it undeliverable; missing endpoint makes it unevaluable, which is the failure mode of gatekeeper programmes worldwide.

The remaining suicide material in this chapter divides as follows. The national strategy as a governance instrument, its delivery vehicle and its financing are taken up in their own section; the national data series read as a distribution belongs with suicide epidemiology; the occupational category argument belongs with farmer suicide. The mechanisms and the international evidence base for all three levers remain with the Somatic-symptom, sexual, impulse-control disorders and suicide notes, and the family-level task with the Child behavioural, emotional and other disorders notes.

Means restriction is the only one of the three levers with an Indian legal instrument and an Indian national effect estimate; gatekeeper training is the only one a district can commission for itself and has no national coverage figure at all; media guidance binds print alone and is evaluated on the content of reporting, never on deaths.

17 · Epidemiology of suicide in India and high-risk groups

Every national suicide figure you will ever quote is the output of a data system, not a measurement of the world. That is not scepticism, it is the first thing a public-health reader has to understand about this topic, and it is what separates an answer that recites a number from an answer that can defend one. India's national suicide statistics come from a police-reporting chain, are published annually, and describe a case category with a legal existence. Everything that follows about counts, rates, trends and high-risk groups is conditioned by that fact.

The clinical material adjacent to this section is owned elsewhere and is not restated. Individual risk assessment, the weighing of risk against protective factors, structured evaluation and the limits of prediction belong to the Somatic-symptom, sexual, impulse-control disorders and suicide notes, as do the aetiological frameworks that explain why populations differ. Custodial suicide prevention, including screening on reception into prison and the advisory machinery around it, belongs to the Forensic ethics, consent and legal procedure notes. What this section owns is the surveillance instrument itself, the distribution it produces, and what that distribution licenses a programme to do.

17.1 What the national statistic is a statistic of

The published national series is compiled by the **National Crime Records Bureau, the NCRB**, from **police-recorded suicide cases** reported upward by states and union territories and collated for a calendar year. Three consequences follow immediately and each is examinable.

- The unit of count is a *recorded case*, so anything that changes recording changes the series. A family's reluctance to have a death recorded as suicide, a local practice of recording an ambiguous death as accidental, or a change in police procedure will all move the number without anything moving in the population.
- The determination is administrative rather than clinical. No psychological autopsy stands behind a national count, and the intent judgment is made in a legal rather than a diagnostic frame.
- The compilation is annual and closed, which means the whole apparatus is retrospective by at least a full year. A surveillance system with that latency describes the past accurately and cannot warn about the present.

RULE **Name the recording system before you quote its number.** "The suicide rate in India is X" is an incomplete sentence. The complete sentence is "the police-recorded suicide rate in India for the year Y, per lakh of the projected mid-year population, is X". Everything a candidate is likely to get wrong about this topic is contained in the words the short version leaves out.

17.2 The count and the rate, and how they differ

For the most recent published year the Bureau **recorded 1,70,746 suicides in India in 2024, against 1,71,418 in 2023**. The corresponding rate is defined as the incidence of suicides per one lakh of population, a lakh being 1,00,000 people, and the recorded rate is a series rather than a constant, so a figure quoted for an earlier reference year in the Somatic-symptom, sexual, impulse-control disorders and suicide notes belongs to an earlier point on the same series and does not compete with this one: for **2024 the rate is 12.2 per lakh, against 12.3 in 2023**.

A count and a rate answer different questions and are not interchangeable. The count answers a service-planning question, because a mortuary, a crisis service and a bereavement service all deal in people rather than in rates. The rate answers a comparison question, because only a rate can be set beside another year, another state or another country. A year in which the count falls while the population grows will show a falling rate for two different reasons at once, and separating them requires the denominator.

1,70,746	12.2	Recorded case	Over a year
POLICE-RECORDED SUICIDE CASES, INDIA, 2024	PER LAKH OF THE PROJECTED MID-YEAR POPULATION, 2024	THE UNIT BOTH FIGURES COUNT, NOT A CLINICAL DETERMINATION	LATENCY BEFORE A CLOSED CALENDAR YEAR IS PUBLISHED

17.3 The denominator, which is where rates are quietly built

India conducts a census at long intervals, so for every year that is not a census year the population **denominator** is not counted but projected. The published rate for the most recent year uses a projected mid-year population of 14049.1 lakh, and the rate for any non-census year is computed on a projection of that kind rather than on an enumeration.

That has two implications a resident should be able to state. The first is that a national rate carries the uncertainty of the projection as well as the uncertainty of the numerator, and neither is usually reported with an interval. The second is that projections drift: the further a year sits

from the last enumeration, the more the denominator depends on assumptions about fertility, mortality and migration, and a small systematic error in the denominator moves every rate in the series in the same direction. A trend computed across a long inter-census interval is therefore a weaker object than it looks, and this is worth saying explicitly in any answer that discusses whether the Indian rate is rising.

TRAP **Comparing a rate for one population against a count for another.** The classic version is to set a national rate beside a group's absolute number and conclude that the group is over-represented. A count of deaths in any group is a function of the size of the group, and without the group's own denominator it says nothing about risk. Where a valid denominator for a subgroup does not exist, the honest position is that the group's rate cannot be computed, not that its count is small or large.

17.4 What the system cannot see, and what else could see it

Undercount in a police-based series is not a defect to be apologised for; it is a structural property with a describable mechanism. Deaths that never reach police attention are absent; deaths that reach it but are recorded under another manner are misclassified; deaths recorded correctly but not transmitted upward are lost. Each of those acts in one direction, so the national series is best understood as a floor rather than an estimate with symmetric error.

The direction of that structural bias has a practical corollary that examiners like. If the series is a floor, then a district or a state reporting a low recorded rate has two possible explanations, a genuinely low incidence and a weak recording chain, and the second explanation is at least as common as the first. A programme that reads a low recorded rate as reassurance and withdraws attention from that district has therefore drawn exactly the wrong conclusion from a weak signal. The correct reading of an unexpectedly low figure is that the recording chain should be examined first and the population second.

India also maintains civil registration of deaths and a sample-based demographic survey, and the natural analytic move is to compare a police-derived suicide count against a registration-derived one for the same period. That comparison is the correct way to ask how large the undercount is. This chapter's sources do not establish the size or the direction of the difference for India, so no reconciliation figure is offered here, and a candidate should say that the comparison is the right method and that its published answer is what they would look for.

17.5 High-risk groups, read as population strata

A **high-risk group** in a public-health frame is not the same object as a high-risk patient. It is a stratum of the population whose rate exceeds the population rate, identified so that a finite intervention can be aimed somewhere rather than everywhere. Three things have to be true before a stratum is useful for targeting, and the national series satisfies only the first of them for most groups.

- **The stratum must be identifiable in the data.** The national record carries profile characteristics collected on the police form, so the series can be cut by sex, by age band, by state and by occupation and other social descriptors. The definitions of those recorded

categories, and what they will and will not bear, are the argument taken up where farmer suicide is taught in this chapter, and they are not repeated here.

- **The stratum must have its own denominator.** Without a population count for the group itself, a group rate cannot be computed and over-representation cannot be established. This is the point at which most confidently quoted group statistics fail.
- **The stratum must be reachable.** A group that can be described but not addressed through any existing service or regulatory route generates a finding, not a programme.

The third condition is worth dwelling on because it is where epidemiology becomes service planning. Strata differ enormously in how reachable they are. A group defined by an occupation can be reached through the sector that employs it and through the regulator that governs its inputs. A group defined by contact with a health service can be reached at that contact, which is why service-contact populations attract prevention effort out of proportion to their share of deaths. A group defined only by a demographic characteristic, such as an age band, has no institution attached to it at all, and reaching it requires a mass channel such as schooling, broadcast media or a helpline. When you are asked to design a targeted intervention, name the institution before you name the intervention, because the institution is what makes targeting possible.

Groups conventionally discussed in the Indian literature include young adults, students, people in agrarian occupations, people in custody and people with severe mental illness in contact with services. Naming a group is legitimate; attaching a proportion to it is not, unless the proportion carries its own denominator. This section prints no group-specific proportions, because none is established in its sources, and that absence should be stated in an answer rather than filled from a slide.

PITFALL

Confusing relative and absolute risk when planning. A stratum can carry a high relative risk and contribute few deaths because it is small, while a low-risk but very large stratum contributes most of the total. Targeting the first reduces risk for individuals; only action on the second moves the national count. A plan that names only high-relative-risk groups will be defensible clinically and useless at population scale, and this is the single most common weakness in resident-written prevention plans.

17.6 Inter-state variation, and why it is the useful unit

Health is administered by states, so the national rate is an average over administrative units that differ in almost everything that matters: recording practice, service availability, urbanisation, the dominant method available, and the age structure of the population. A national figure is therefore the least actionable number in the series. The state figure is the one a programme can be built on, and the district figure is the one a service can be built on.

Two cautions apply when reading variation between states. Differences in the recorded rate confound true differences in incidence with differences in recording, and there is no way to

separate them from the published series alone. And states differ in age structure, so a raw comparison between them is not a like-for-like comparison; the standard remedy is **age standardisation**, and any inter-state comparison that has not been standardised should be described as a comparison of **crude rates**. This chapter prints no state figures, because none is in its sources.

17.7 Reading the trend as a data series

The two most recent published years sit close together, with the count and the rate both marginally lower in the later year. A one-year movement of that size is not a trend and should never be described as one. A movement in a routinely collected series can arise from a real change in incidence, from a change in recording behaviour, from a change in the projected denominator, or from ordinary year-to-year variation, and a single comparison cannot distinguish them.

What an examiner is looking for is the discipline to say all four aloud, in that order, before offering an interpretation. If a candidate must commit, the defensible statement is that the national recorded rate has been broadly stable across the two most recent published years, that the level is what matters for planning rather than the direction of the last step, and that any claim about a sustained trend requires the series over many years together with a stated treatment of the denominator.

IN INDIA

The operational decision this section drives is what you plan against. Because the national series is retrospective by more than a year, is administratively defined, and averages over states that differ in recording as much as in incidence, a district service should not wait for the national publication to tell it what its problem is. It should treat the national figure as the background level and build its own local numerator from the sources it can see directly, meaning its own casualty and mortuary records and its own crisis presentations, while accepting that a locally built count is a count and never a rate unless the district's population denominator is stated with it.

17.8 What the distribution licenses, and what it does not

The final discipline is to be clear about what an epidemiological description entitles anyone to do. It licenses the aiming of finite resources at strata with a demonstrably higher rate, provided that the rate has a denominator. It licenses surveillance of the strata over time, which is the only way to know whether targeting worked. It does not license a prediction about any individual, and it does not license the inference that a correlate observed at population level operates as a cause in a person.

QUANTITY	WHAT IT ANSWERS	WHAT IT NEEDS TO BE VALID	WHAT IT CANNOT DO
National recorded count	How many recorded cases arose in a closed calendar year	A stated recording system and a stated year	Support comparison across years or places without a denominator
National recorded rate	How the level compares with another year, state or country	The same numerator plus a stated population base, projected for non-census years	Describe any subgroup, and cannot be applied to an individual
Group rate	Whether a stratum is over-represented	A numerator and a denominator for that same group	Be computed at all where the group's population base is unavailable
Trend	Whether the level is moving	Many years, a stable case definition and a stated denominator method	Be read from a single year-on-year step

India's national suicide figures are police-recorded counts and rates for a closed calendar year, computed for non-census years on a projected mid-year population, so they describe the past accurately, warn about nothing in the present, and support a statement about any subgroup only where that subgroup has a denominator of its own.

Stigma and anti-stigma interventions

18 · Stigma: concepts, types (public, self, structural) and measurement instruments

Stigma is not one thing, and a service that treats it as one thing will intervene in the wrong place. The word covers at least four distinct objects, each sitting in a different location, each moved by a different intervention and each measured by a different instrument. Typing stigma is therefore not an academic refinement; it is the step that decides whether an anti-stigma budget buys a media campaign, a workforce training programme, a change to an insurance rule or a support group for families. This section supplies the typology, its consequences for access to care, and the measurement layer that lets any of it be evaluated.

Two neighbouring accounts are owned elsewhere and are named rather than rebuilt. The social-psychological mechanism behind stigma, prejudice as a structure of belief, feeling and behaviour, the theoretical lineage that explains how in-groups and out-groups form, and the intergroup evidence base, belongs to the Social, Learning and Personality Psychology notes. The clinical reading, in which stigma is treated as a variable inside one consultation and divided into what a patient has met, what a patient anticipates and what a patient has internalised, with concealment as the behavioural pathway to a worse outcome, belongs to the Consultation-liaison, geriatric and transcultural psychiatry notes. What follows is the population and structural reading, which no other chapter carries.

18.1 Why a service types stigma rather than deploring it

Deploring stigma is free and changes nothing. A public-health reading asks four questions of any stigma phenomenon, and the answers differ by type.

- **Where does it live?** In a population, in a person, in a rule, or in a family.
- **Who holds it?** The answer identifies the respondent any later measurement will have to use.
- **What does it do?** Specifically, what does it do to whether a person seeks care, receives care and keeps a role, because a stigma that changes nothing is a sentiment rather than a public-health problem.
- **What would count as evidence that it had shifted?** This has to be answerable before the intervention is designed, not afterwards.

Running those four questions produces the typology below rather than presupposing it, which is the right order and the one an examiner rewards. Note that the fourth question is the one that most often has no good answer, and that a type of stigma for which no endpoint can be named is a type nobody will be able to show they reduced.

RULE **Locate the stigma before you choose the intervention.** An intervention aimed at the wrong location will produce a measurable change in the wrong outcome and no change in access. A campaign that improves public attitudes will not amend a discriminatory insurance rule, and a support group that reduces a patient's shame will not alter what a stranger believes. Both may be worth doing; neither substitutes for the other.

18.2 Public stigma

Public stigma is what a population believes, feels and does about people with mental illness. It is held by other people, it is distributed across a society and it is the type most people mean when they use the word without qualification. Its behavioural expression is **discrimination**, which is the part that has consequences: a tenancy refused, a job offer withdrawn, a marriage negotiation abandoned, a family member excluded from a decision.

Two distinctions matter for measurement, and both are frequently collapsed. The first is between an attitude and a behaviour: what someone reports believing about mental illness in the abstract predicts, but does not equal, what they do when a person with mental illness applies for a room. The second is between the general public and specific gatekeeping populations, notably employers, landlords, teachers, police and health professionals. Health professionals are a gatekeeping population in their own right, and stigma held inside the health system is the version with the most direct effect on whether a person gets treated at all.

18.3 Self-stigma

Self-stigma is what happens when a person who is aware of a public stereotype applies it to themselves and endorses it. The location has moved: this now lives inside the person, and its consequences are demoralisation, withdrawal from roles the person could still occupy, and reduced help-seeking. Its practical importance in a service is that it is a mechanism by which a

public phenomenon becomes a clinical one, and that it is largely invisible to any measure of public attitudes.

The service consequence is specific. Self-stigma degrades exactly the behaviours a chronic-illness service depends on: turning up, disclosing symptoms, accepting a rehabilitative role, applying for entitlements. A programme that reduces public stigma may leave self-stigma untouched, because a person can know that attitudes in general have improved and still hold the belief about themselves, so the two require separate measurement and separate intervention.

18.4 Structural stigma

Structural stigma is the version that requires no person to hold any attitude at all. It is the accumulation of rules, budgets, laws, institutional practices and physical arrangements that disadvantage people with mental illness, and it persists after every individual in the system has changed their mind. Recognising it is what turns a clinician into a public-health practitioner, because it is the only type of stigma whose remedy is administrative rather than educational.

Its examples in any health system are structural rather than attitudinal: mental health services funded at a lower share than the burden they carry, insurance products that exclude or limit psychiatric cover, facilities sited away from general hospitals so that attending one is itself a disclosure, records systems that make a psychiatric contact more visible than a medical one, and application forms that ask about psychiatric history in terms no medical condition attracts. Each of those is fixable by a decision, and none is fixable by a campaign.

TRAP **Answering a structural stigma question with an attitude intervention.** Asked how to reduce stigma affecting access to care, most candidates describe awareness work. Awareness addresses public stigma. Where the barrier is a funding rule, a siting decision or an insurance exclusion, the intervention is a change to that rule, and naming the rule is what distinguishes a public-health answer from a well-meaning one.

18.5 Associative stigma

Associative stigma, also called courtesy stigma, is the discredit that attaches to people connected to the affected person rather than to the person themselves. It falls on family members, on carers, and in some settings on the professionals who work in the field. It is the type most under-measured and, in a family-based care system, arguably the most consequential, because it operates on the people whom services rely on to deliver most of the care.

Its behavioural signature is concealment at the family level rather than at the individual level, with predictable effects on presentation delay, on treatment continuity and on marriage and employment decisions taken about other members of the household. A service that asks only the patient about stigma will not detect it.

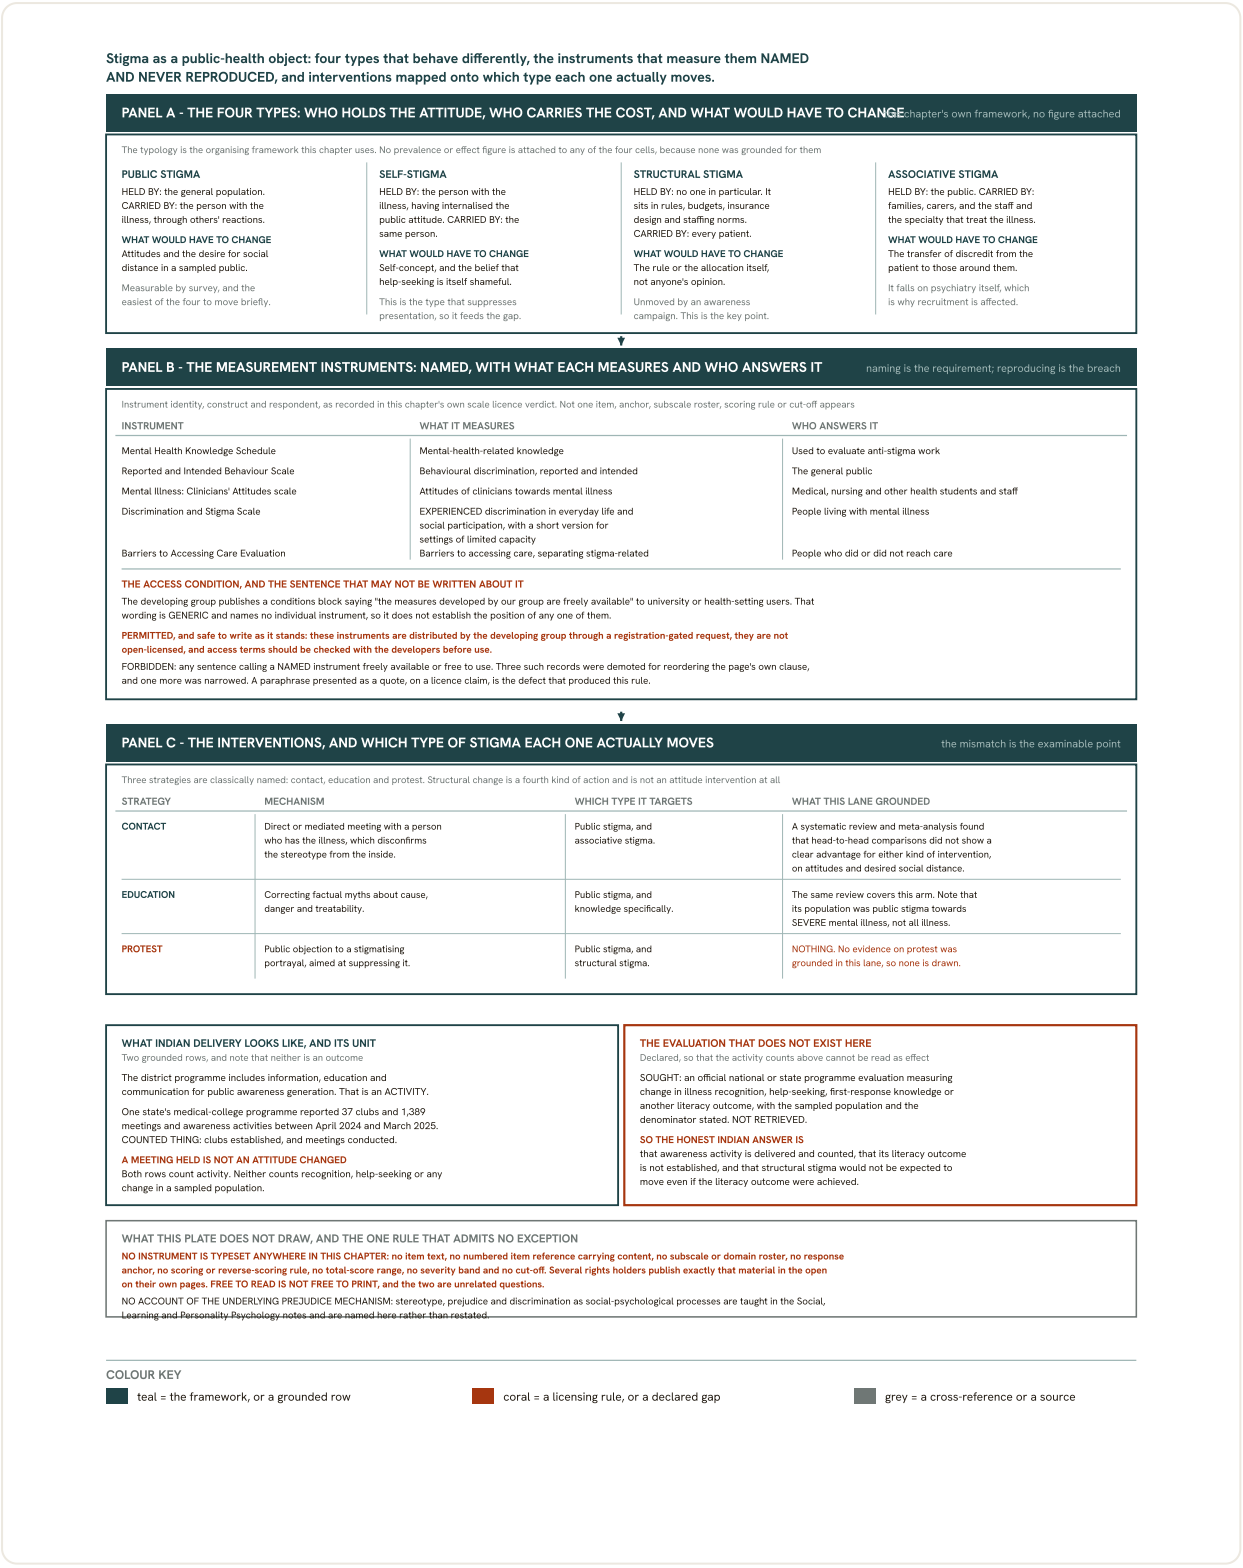


Fig 39.7 Stigma typed, measured and intervened on, with the measurement instruments named and none reproduced. The four types differ in who holds the attitude and who carries the cost, and structural stigma is the one an awareness campaign would not be expected to move. The instruments are named with their construct and respondent only. On access, the developing group's conditions block says its measures are freely available to university or health-setting users, but the wording is generic and names no individual instrument, so no named instrument may be called freely available; records that did so were demoted for reordering the page's own clause. On interventions, head-to-head comparisons showed no clear advantage for contact over education, and nothing on protest was grounded. Indian delivery is counted as activity, awareness generation through the district programme and 37 clubs with 1,389 meetings in one state between April 2024 and March 2025, while no official evaluation of any literacy outcome was retrieved.

MNEMONIC

Them, Me, The rules, Mine

The four locations of stigma, in the order in which a service should look for them. What others hold, what the person has taken on, what the system enforces without holding any view at all, and what falls on the family. Each location has its own intervention and its own measure, and an answer that covers all four has typed the problem rather than deplored it.

18.6 Where the clinical cut fits

A clinician assessing one patient will more often use a different cut of the same material, dividing stigma into what the person has actually met, what they anticipate meeting, and what they have internalised. That reading is taught in the Consultation-liaison, geriatric and transcultural psychiatry notes and is not re-derived here. What is worth stating in a public-health section is how the two cuts relate: the clinical cut is organised by the patient's relationship to the experience and is therefore useful in a consultation, while the population cut is organised by where the phenomenon lives and is therefore useful when choosing an intervention and an evaluation endpoint. They are not rival taxonomies and a candidate should not be trapped into defending one against the other.

18.7 Measuring stigma: what the instruments are for

Item by item this is the part of the topic candidates find hardest, because the instrument names are numerous and superficially interchangeable. They are not. Each was built to measure a specific construct, in a specific respondent, for a specific purpose, and using one in place of another produces a number that answers a question nobody asked. This section names them and describes them; it prints no item, no domain roster, no response option, no scoring rule and no threshold from any of them, and neither should you.

INSTRUMENT	CONSTRUCT IT MEASURES	WHO COMPLETES IT	WHAT IT CANNOT DO
Mental Health Knowledge Schedule (MAKS)	Mental health related knowledge in the general population	Members of the public	Detect discrimination, since knowing something is not doing something
Reported and Intended Behaviour Scale (RIBS)	Reported past and intended future contact behaviour toward people with mental illness	Members of the public	Establish that intended behaviour was enacted
Discrimination and Stigma Scale (DISC), and its ultra-short form (DISCUS)	Experienced discrimination across areas of everyday life and social participation	People with lived experience of mental illness	Measure what the public believes, since it records the receiving end only
Mental Illness: Clinicians' Attitudes scale (MICA)	Attitudes toward mental illness and psychiatry within a health workforce	Health students and health professionals	Stand in for public attitudes, since the respondent group is a gatekeeping profession
Internalized Stigma of Mental Illness scale (ISMI)	Internalised stigma in the affected person	People with lived experience of mental illness	Say anything about the surrounding population
Attribution Questionnaire (AQ) and Self-Stigma of Mental Illness Scale (SSMIS)	Attributional reactions to mental illness, and the self-directed form of the same process	Members of the public, and affected individuals respectively	Substitute for one another, despite sharing a theoretical origin
Community Attitudes toward Mental Illness (CAMI) and the Perceived Devaluation-Discrimination scale (PDD)	Community attitudes, and perceived societal devaluation as seen by the affected person	Community samples, and affected individuals respectively	Be read as measures of the same thing, since one records what a community holds and the other what a person believes a community holds
Barriers to Accessing Care Evaluation (BACE)	Barriers to accessing mental health care, distinguishing stigma-related barriers from other kinds	People who have needed or sought care	Attribute a barrier to a specific policy or institution

Two structural points about that table are worth more than the names in it. First, the respondent column is doing most of the work: some of these instruments interrogate the population and some interrogate the person receiving the population's behaviour, and a study that mixes them has measured two different phenomena and averaged them. Second, the last column is the one to reproduce in an answer, because knowing what an instrument cannot do is the evidence that you know what it does. Social distance measures, a family of instruments in this literature rather than a single named tool, sit alongside these and are usually treated as a behavioural proxy for public attitude.

GOOD TO REMEMBER

If you are going to evaluate anything you do about stigma, choose the respondent before you choose the instrument. If your intervention acts on staff, measure staff. If it acts on the public, measure the public. If your claim is that access improved, then the only respondent who can confirm it is the person who did or did not access care, and the instrument has to be one that asks them about barriers rather than one that asks anybody about attitudes.

18.8 Access to the instruments, and what is actually established about it

This is where a chapter can quietly go wrong, and it is worth being exact. The research group that developed a large part of the stigma instrument suite publishes a conditions block on each instrument's page. That block states that **the measures developed by the group are freely available to users who work or study in a university or health setting**, and it also forbids modification, forbids passing the measures to a third party and forbids charging for them. Every sentence of that block speaks of "the measures" generically and names no individual instrument.

The consequence is precise and it is the teaching point. Because the permissive sentence names no instrument, it does not establish the access terms of any particular one. Whether the clinicians' attitudes scales in that suite, the reported and intended behaviour measure, or the discrimination scale and its short form are individually free to a university or health-setting user is *not established* by the text that was sought and read. The defensible formulation, and the one to use, is that these instruments are distributed by the developing group through a registration-gated request, are not open-licensed, and that access terms should be confirmed with the developers before use. Other instruments in the field sit with individual rights holders or with commercial publishers and carry their own, different terms.

PITFALL

Assuming that a scale you can read online is a scale you may reproduce. Several rights holders publish item text, response options and scoring information openly on their own pages, for their own good reasons. Free to read is not free to print, not free to translate, not free to embed in a records system and not free to modify. A dissertation, a poster or a teaching handout that reproduces an instrument without checking its terms is a copyright problem regardless of how easy the material was to find.

18.9 Choosing an endpoint, which is the whole examination question

The last step is to connect type, instrument and endpoint into one line of reasoning, because that is what a viva question about stigma is really testing. If the target is public stigma, the endpoint is population knowledge, attitude or reported behaviour, and the measure is a public-facing instrument. If the target is the health workforce, the endpoint is professional attitude and the respondent is staff. If the target is self-stigma, the endpoint is internalised stigma in affected people. If the target is structural stigma, no attitude instrument will do at all, and the endpoint is the rule, the budget line or the entitlement itself, measured as a document rather than as a questionnaire. And if the claim is that any of this improved access, the endpoint is a barrier measure completed by people who needed care.

Anti-stigma interventions and their comparative evidence are taken up in their own section of this chapter, and mental health literacy, which overlaps with public knowledge but is a distinct programme class with its own outcomes, in the section after it. Neither of those names an instrument; the measurement layer stays here.

Location	Respondent	Rules	Not established
WHAT TYPES STIGMA: PUBLIC, SELF, STRUCTURAL, ASSOCIATIVE	WHAT DECIDES WHICH INSTRUMENT IS VALID	THE ONLY EVIDENCE THAT STRUCTURAL STIGMA MOVED	PER-INSTRUMENT ACCESS TERMS IN THIS SUITE

Type the stigma by where it lives, choose the respondent before the instrument, and remember that structural stigma is the only type whose remedy is a change to a rule and whose evidence is a document rather than a questionnaire.

19 · Anti-stigma interventions: contact, education and protest strategies and evidence

Anti-stigma work is a programme class, not a good intention. Three strategies are conventionally named, contact, education and protest, and a fourth family of structural measures sits alongside them without usually being counted as a strategy at all. Each has a design, a cost, a delivery workforce, a plausible mechanism, a characteristic failure mode and an outcome that can in principle be measured. This section treats them in exactly those terms, because a resident who can say what an anti-stigma programme would look like on the ground and how it would be evaluated has answered a public-health question, while a resident who can only assert that stigma should be reduced has not.

The social-psychological mechanism behind the contact strategy, the conditions under which intergroup contact reduces prejudice and the meta-analytic evidence establishing it, belongs to the Social, Learning and Personality Psychology notes and is named here rather than rebuilt. The stigma typology on which everything below depends, and the measurement instruments used to evaluate any of it, are taught in the section of this chapter that owns them; no instrument is named here.

19.1 Four families, and what each is actually trying to change

Before comparing the strategies it is worth being explicit that they act on different things, because most of the confusion in this topic comes from treating them as interchangeable ways of doing the same job.

- **Contact** puts members of a target audience into structured interaction with people who have lived experience of mental illness. Its target is the audience's beliefs and, more importantly, its behavioural intentions.
- **Education** supplies correct information in place of a false belief. Its target is knowledge, and its assumption is that stigmatising belief is partly an information deficit.
- **Protest** publicly objects to a stigmatising representation or practice and demands its withdrawal. Its target is not the audience at all, but the producer of the representation.

- **Structural measures** change a rule, a budget, a siting decision or an entitlement. Their target is the disadvantage itself, and they require no one to change their mind.

Two things follow immediately. First, only the last of the four addresses structural stigma, so a programme built entirely from the first three has left the most durable form of the problem untouched. Second, protest is aimed at a different recipient from the other three, which is why its evaluation looks different and why it carries a risk the others do not.

RULE **Match the strategy to the type of stigma and to the audience that holds it.** Contact and education act on people who hold attitudes. Protest acts on people who produce representations. Structural measures act on institutions. An intervention aimed at the wrong recipient can be delivered perfectly and still be incapable of changing the outcome it claims.

19.2 Contact as a programme rather than as a mechanism

Taking the mechanism as read, the programme questions are the ones a service actually has to answer. Who are the people with lived experience who will deliver it, how are they recruited, prepared and supported, and what happens to them afterwards? Contact-based programmes run on the disclosure of individuals, which is a genuine cost borne by the very people the programme exists to help, and a programme design that does not address the safeguarding, remuneration and sustainability of its speakers is exploiting them.

The other programme decisions are the audience, the dose and the medium. Audience matters because a gatekeeping group such as health staff, employers, landlords, teachers or police will produce a larger change in real-world discrimination per person reached than the general public will. **Dose**, meaning the number and spacing of encounters, matters because a single session is the commonest design and the least likely to produce durable change. Medium matters because live and mediated delivery differ in reach and in cost in opposite directions, and a filmed or written encounter can reach a population that no live programme could, at the price of some of what makes the encounter work.

19.3 Education as a programme

Educational anti-stigma work replaces a false belief with a correct one, and its natural home is any setting with a captive audience and an existing teaching structure: schools, colleges, professional training and workplace induction. Its advantages are that it is cheap, scalable, deliverable by people who are already employed to teach, and easy to standardise. Its limitations are equally clear. Knowledge and attitude are different constructs and correcting the first does not reliably move the second; attitude and behaviour are different again; and a message pitched to correct one stereotype can, if it is framed carelessly, reinforce the belief that the condition is a fixed and alien state.

The framing question inside educational work is genuinely contested and is worth knowing. Messages that present mental illness as an illness like any other are intended to reduce blame, and they can succeed at that while making the audience perceive greater difference and greater unpredictability. Messages that emphasise continuity between ordinary distress and disorder reduce perceived difference but can dilute the case for treatment. There is no framing that

optimises everything, and an examiner will be satisfied by a candidate who can name the trade-off rather than one who asserts a winner.

19.4 Protest, and why it is the least taught of the three

Protest challenges a specific stigmatising representation: an advertisement, a film, a headline, a product name, a public statement, an institutional practice. It is the only strategy of the three whose outcome is immediate, visible and binary, because either the material is withdrawn or it is not.

Its limitation is that suppressing a representation is not the same as changing the belief behind it, and there is a well-recognised risk of **reactance**, a defensive or contrary response in an audience that reads a public demand for withdrawal as censorship. It is therefore usually described as effective at the level of the representation and unreliable at the level of the attitude. In practice it is best understood as a tool for removing amplifiers rather than for persuading anyone, and it sits most comfortably alongside the structural family, because both act on institutions.

TRAP **Claiming attitude change from a protest outcome.** A withdrawn advertisement is evidence that a producer responded, not that a population's beliefs moved. The two are frequently reported together in programme write-ups, and separating them is exactly the discrimination an examiner is looking for.

19.5 Structural measures, which are where the durable gains are

Structural measures in anti-stigma work act on the rules. Removing an exclusion of mental illness from an insurance product, equalising an entitlement, siting a psychiatric service inside a general hospital, removing a psychiatric history question from a form that asks nothing comparable about physical illness, and correcting a funding allocation that is out of proportion to burden are all anti-stigma interventions, and none of them is a campaign.

Three features make them attractive and one makes them hard. They are durable, because a rule stays changed after enthusiasm fades; they are universal, because they apply to everyone the rule touches rather than to those who attended; and their evidence is a document, so evaluation is cheap and unambiguous. What makes them hard is that they require access to decision-makers rather than to audiences, which is a different competence and a longer timescale, and it is the competence a psychiatrist in a public system is uniquely placed to supply.

19.6 What the comparative evidence actually says

The question candidates most want answered is whether contact beats education. The best available synthesis is more useful and less satisfying than a winner: in a systematic review and meta-analysis of interventions to reduce stigma towards people with severe mental illness, **head-to-head comparisons did not show a clear advantage for either kind of intervention**. The outcomes in view were stigmatising attitudes and the desire for social distance, and the population in view was public stigma towards people with severe mental illness, defined in that review as schizophrenia, psychosis or bipolar disorder.

Read the scope before you use the conclusion. It concerns public stigma towards severe mental illness, not stigma towards common mental disorders, not self-stigma, not structural stigma and not the health workforce. It concerns short-term attitudinal and social-distance outcomes, not discrimination as experienced by patients and not access to care. A finding of no clear head-to-head advantage within that scope is not a finding that the two strategies are equivalent in every setting, and it is certainly not a finding that either changes behaviour.

The practical implication is liberating rather than deflating. If neither strategy is clearly superior on attitudinal outcomes, then the choice between them should be made on the grounds a programme manager can actually control: cost per person reached, availability of a willing and supported speaker group, existing delivery structures, the ability to reach the audience whose behaviour matters most, and the feasibility of measuring anything afterwards.

PITFALL

Reading "no clear advantage" as "both work well". A head-to-head comparison that fails to separate two interventions says nothing on its own about whether either produced a meaningful change against no intervention, and nothing at all about whether any change persisted or reached behaviour. Quote the comparison for what it is, a comparison, and keep the separate questions of efficacy, durability and behavioural transfer separate.

19.7 Targeting: everybody, or the people who decide

The last design decision, and the one with the largest effect on what a fixed budget buys, is whom to aim at. A whole-population campaign has the appeal of universality and the disadvantage that most of the people it reaches will never be in a position to discriminate against anyone with mental illness. A targeted programme aimed at a gatekeeping group accepts a smaller reach in exchange for a much higher density of consequential decisions per person reached, because every member of that group is repeatedly in a position to grant or withhold something: a job, a tenancy, a school place, an arrest decision, a referral, a bed.

Health professionals deserve a separate mention, because they are the gatekeeping group a psychiatric service can reach at zero marginal cost and the one whose stigma has the most immediate effect on whether a person is treated. Diagnostic overshadowing, in which a physical complaint in a person with known mental illness is attributed to the illness and not investigated, is the concrete form the problem takes inside a hospital, and it is a discrimination outcome rather than an attitude, which means it can be audited from records rather than surveyed. An anti-stigma programme aimed at one's own institution with an auditable endpoint of that kind is more evaluable than most of what appears in the published literature.

19.8 Endpoints, decay and the transfer problem

Three measurement problems recur across every anti-stigma programme and a good answer names all three.

The first is **endpoint selection**. Knowledge, attitude, intended behaviour, reported behaviour, experienced discrimination and access to care form a ladder, and each rung is harder to move and

more meaningful than the one below. Almost all published anti-stigma evaluation sits on the lower rungs, and the respondent has to change as you climb: the top of the ladder can only be reported by people who needed care.

The second is **decay**. Attitudinal change after a brief intervention tends to fade, so a design with a single post-intervention measurement taken immediately afterwards cannot distinguish a durable effect from a transient one. A follow-up measurement at a meaningful interval is not a refinement, it is the difference between an evaluation and a satisfaction survey.

The third is **transfer**. Attitude measured in a classroom does not automatically become behaviour in a rental market or a recruitment decision, and the gap between the two is where anti-stigma programmes most often fail without knowing it. Where a programme's real claim is about behaviour, the design has to include a behavioural or access-based endpoint, and if it cannot, the claim has to be scaled back to what was actually measured.

STRATEGY	STIGMA TYPE IT CAN MOVE	WHO RECEIVES IT	ENDPOINT THAT WOULD SHOW IT WORKED	CHARACTERISTIC FAILURE
Contact	Public, and workforce attitudes	An audience, ideally a gatekeeping one	Attitude and intended behaviour in that audience, remeasured after an interval	Single-session dose, unsupported speakers, and no follow-up measurement
Education	Public knowledge, and workforce knowledge	A captive audience in an existing teaching structure	Knowledge in that audience, and separately any attitude change	Knowledge improves while attitude does not, and the two are reported as one
Protest	The representation, not the belief	The producer of a stigmatising representation	Withdrawal or amendment of the material	Reactance in the audience, and attitude change claimed from a withdrawal
Structural measures	Structural stigma	An institution or a policymaker	The amended rule, budget line or entitlement, read as a document	Never attempted, because it needs access to decision-makers rather than to audiences

GOOD TO REMEMBER

If you are asked to reduce stigma in your own hospital, start with the workforce and with the rules, because those are the two things you can actually change and both act directly on whether a patient is treated well. Contact and education aimed at your own staff are deliverable inside existing teaching time, and the structural audit costs nothing: walk the pathway a psychiatric patient takes through your institution and write down every point at which it differs from the pathway a medical patient takes for no clinical reason. That list is an anti-stigma intervention plan.

MNEMONIC

Meet, Teach, Object, Amend

The four families in the order of increasing durability. Meeting someone changes an audience, teaching changes what an audience knows, objecting removes a representation, and amending a rule changes the disadvantage itself and stays changed. Most programmes stop at the first two, and most of the durable gains are in the last.

Audience	Producer	Institution	No clear advantage
WHAT CONTACT AND EDUCATION ACT ON	WHAT PROTEST ACTS ON	WHAT A STRUCTURAL MEASURE ACTS ON	CONTACT AGAINST EDUCATION, HEAD TO HEAD, PUBLIC STIGMA IN SEVERE MENTAL ILLNESS

Contact and education have not been separated head to head on public stigma towards severe mental illness, so choose between them on cost, audience and feasibility, and remember that the only strategy that reliably outlives the enthusiasm of the people running it is the one that changes a rule.

20 · Mental health literacy and community awareness programmes

A treatment gap is closed one step at a time, and the first step is inside somebody's head. Before a person reaches a service they have to notice that something is wrong, decide that it is the kind of wrong a health system deals with, believe that something can be done about it, and know where to go. **Mental health literacy** is the name for that stock of knowledge and belief in a population, and awareness programmes are the deliberate attempt to raise it. This is a programme class in its own right, with its own audiences, its own delivery vehicles and its own outcomes, and it is not a synonym for anti-stigma work.

The distinction matters because the two are constantly conflated. Anti-stigma work targets attitude and discrimination and is taken up in its own section of this chapter. Literacy work targets recognition, help-seeking and first-response knowledge. They overlap, because a population that believes mental illness is a moral failing will not seek help for it, but they are separable, and a programme that measures one while claiming the other has produced an evaluation that answers the wrong question. No measurement instrument is named in this section; the instrument layer belongs to the section of this chapter that owns stigma measurement.

20.1 What literacy has to contain, derived from the pathway rather than asserted

Rather than reciting a list, it is more useful to derive the components from the pathway a person actually travels, because then each component has an owner and an intervention attached to it.

- **Recognition.** The person, or somebody near them, notices a change and names it as a change. Without this step nothing else happens, and this is the step at which a condition that presents as bodily complaint or as behaviour rather than as reported mood is most often missed.

- **Attribution.** The change is attributed to a health problem rather than to character, misfortune, external agency or the ordinary difficulty of life. This is the step at which explanatory models operate, and it is the step most sensitive to what a community believes about causes.
- **Efficacy belief.** The person believes something can be done. A population that recognises depression and believes it is untreatable will not convert recognition into help-seeking.
- **Route knowledge.** The person knows a specific place to go and what will happen there. "Seek help" is not route knowledge; a named service with a location, a cost and a plausible account of what the first visit involves is.
- **First-response knowledge.** The people around the person know what to do and what not to do in the interval before care is reached, which in this field includes knowing how to respond to talk of self-harm.

Reading the components this way has a practical payoff. It shows that a campaign can raise recognition and change nothing, because the bottleneck was efficacy belief or route knowledge, and it tells a programme designer to find out which step their population is actually stuck at before choosing a message.

RULE **Literacy is only useful up to the capacity of the service it points at.** A successful awareness campaign generates demand. If the district service cannot absorb it, the campaign converts an untreated population into a disappointed one, and the second is harder to reach the next time. Awareness and service capacity are commissioned together, or the awareness work should wait.

20.2 The Indian delivery vehicle that already exists

India does not need to invent a channel for this work, because the district programme already carries one. **Information, education and communication for public awareness generation is a named component of the District Mental Health Programme**, sitting alongside its service, training and other activities. That single fact is worth more in an examination than any number, because it locates the responsibility: public awareness is a programme deliverable with an owner in every district where the programme runs, not an optional add-on that depends on an enthusiastic individual.

Two consequences follow for a resident planning district work. The first is that the mandate already exists, so the question is not whether to do awareness work but how to do it well and how to show it was done. The second is that a component defined as information, education and communication invites the cheapest possible interpretation, which is the production of material. Producing a poster is an activity; changing what a population knows is an outcome; and the distance between the two is the whole difficulty of this topic.

IN INDIA

Because public awareness is a district programme component with a named owner, the decision a resident actually faces is what to do with a mandate that comes with no specified outcome. The defensible course is to choose one bottleneck in the pathway and one audience that can be reached repeatedly through an existing institution, then write down in advance what will be counted and what would count as success. A district that does this produces a small evaluable programme. A district that does not produces material, reports the material, and cannot say at the end of the year whether anybody knows anything they did not know before.

20.3 A worked state example, and what its numbers actually measure

State programmes supply the more instructive examples, because they report what they did. One state programme built around mental health clubs in government medical colleges reported, for a single year of operation, that **37 clubs were established across government medical colleges, conducting 1,389 meetings and awareness activities, between April 2024 and March 2025**.

Now read those two numbers with the discipline this chapter demands. The first counts *clubs established*, which is an institutional structure, and it is a real and useful thing to count because a standing structure inside an institution is what makes a programme survive the departure of the person who started it. The second counts *meetings and awareness activities conducted*, which is an output. Neither counts people, so neither is a reach figure; neither counts knowledge, so neither is a literacy outcome; and neither carries a denominator, so no proportion can be computed from either. What they establish is that a delivery structure existed and was active, which is a necessary condition for an effect and not evidence of one.

The generalisable lesson is that the institution is a good unit for this work. A college, a school, a workplace and a police training centre are all standing institutions with a captive audience, a teaching structure, a calendar and somebody accountable, which is why programmes built on institutions persist and campaigns built on events do not.

37	1,389	Output	Not established
CLUBS ESTABLISHED ACROSS ONE STATE'S GOVERNMENT MEDICAL COLLEGES	MEETINGS AND ACTIVITIES CONDUCTED IN THOSE COLLEGES IN ONE YEAR	WHAT BOTH FIGURES ARE, RATHER THAN REACH OR KNOWLEDGE	CHANGE IN RECOGNITION OR HELP-SEEKING, NATIONALLY OR IN ANY STATE

20.4 Designing a campaign that could work

Five design decisions determine whether an awareness intervention has any chance, and they should be made in this order.

Audience before message. A campaign aimed at everybody is aimed at nobody. The audiences worth choosing are those who are either at high risk, or in repeated contact with people at risk, or in a position to open the door to care. In practice that means adolescents and young adults, frontline health workers, teachers, employers and the police, and each requires a different message.

Bottleneck before content. Establish which pathway step the audience is stuck at. If recognition is intact and route knowledge is absent, a campaign explaining what depression is will change nothing that a card with an address would not have changed better.

Channel by audience, not by fashion. The channel that reaches a college cohort is not the channel that reaches an agricultural community or an older population, and the cheapest channel per contact is rarely the most effective per contact. Where a helpline or a service is named in the message, the channel and the service must be aligned so that the door is open when the message lands.

Language and idiom. A message has to arrive in the language people actually use for distress. Translation of a technical message into a regional language is necessary and not sufficient, because the idiom in which distress is expressed differs from the clinical vocabulary in every language. The cultural material behind this belongs elsewhere in this chapter and in the Consultation-liaison, geriatric and transcultural psychiatry notes; the programme point is simply that a message that does not use the population's own words for suffering will not be recognised as being about them.

Dose and repetition. A single exposure decays. Programmes embedded in a recurring institutional calendar accumulate; single events do not, however well attended they were.

A sixth consideration cuts across all five and is specific to this field. Awareness messages about mental illness can carry unintended harm in a way that messages about, say, hand hygiene cannot: material about suicide can affect a vulnerable reader, material that presents a condition as permanent can depress efficacy belief in exactly the people it hoped to reach, and material that names a diagnosis loosely can teach a population to apply a label rather than to seek an assessment. Every awareness message in this field therefore needs a review step that asks who the most vulnerable reader is and what this text does to them, and that step is cheap, is almost always omitted, and is the sort of thing a psychiatrist attached to a district programme is uniquely qualified to insist on.

PITFALL

Substituting an awareness day for an awareness programme. An annual observance is a communication opportunity, and it is often the only thing a service does. It reaches a self-selected audience once, it has no bottleneck analysis behind it, it produces an attendance figure rather than an outcome, and it recurs annually without accumulating. If a day is all the resource there is, attach it to a standing institutional structure so that it becomes the visible part of something continuous rather than the whole of it.

20.5 Measuring literacy, and what India does not know about itself

Literacy outcomes fall into an ordered set, and knowing the order is most of what an examiner wants. Activities delivered is the weakest, followed by people reached, then knowledge and beliefs in a sampled population, then help-seeking behaviour, then service contact, then outcomes among those who reached care. The rungs above knowledge require the person's behaviour rather

than their report, and the rungs above people reached require a sample and a denominator rather than a tally.

Here the position must be stated honestly. No official national or state programme evaluation measuring a change in illness recognition, help-seeking, first-response knowledge or any other literacy outcome, with the sampled population and denominator stated, was retrieved for India. That was sought through the national health mission, the national institute, and health-department sources for individual states. What is published is what programmes did, not what populations came to know. A resident asked what proportion of Indians can recognise depression, or whether awareness activity has changed help-seeking, should say that no such national figure is established, name what would be needed to establish it, and describe what is published instead.

TRAP **Reporting an activity count as evidence of impact.** Programme reports print what is countable, and what is countable is sessions, posters, participants and events. None of those is a literacy outcome. The correct sentence names the count for what it is, an output, and states separately that the outcome was not measured. Candidates who quote activity totals as achievements are repeating the error the report itself is careful not to make.

There is a second and quieter reason to insist on the ordered set. Programmes are funded on what they can report, so a funder that asks only for activity counts will get a system optimised for activity, and the people running it will not be behaving badly; they will be answering the question they were asked. Changing what is asked for is therefore a more powerful intervention on the quality of awareness work than any exhortation about rigour, and it is a change a clinician sitting on a district or state committee is in a position to propose. Ask for one population-sampled knowledge measure a year, with its denominator stated, and the programme will reorganise itself around producing one.

20.6 How this connects to everything else in the chapter

Literacy work sits upstream of almost every other topic here. It feeds the treatment gap, because recognition and route knowledge are two of the gap's determinants. It feeds suicide prevention, because first-response knowledge in the people around a person at risk is what a gatekeeper programme is trying to install in a workforce, and community awareness is the same intervention aimed at a population. It feeds primary care, because a population that presents earlier presents to a general clinic rather than to a specialist. And it is limited by service capacity, which is why it should never be planned in isolation from the service it points at.

PATHWAY STEP	WHAT THE POPULATION MUST HOLD	INSTITUTION THAT CAN SUPPLY IT	WHAT MEASURING IT REQUIRES
Recognition	The ability to name a change as a change	Schools, colleges, workplaces, frontline health workers	A sampled population and a stated denominator, not an attendance tally
Attribution	The belief that it is a health matter	Community channels in the population's own idiom	The same sample, asked about causes rather than about labels
Efficacy belief	The belief that something can be done	People with lived experience, and clinicians who are visible locally	Belief measured before and after, with an interval long enough to show decay
Route knowledge	A named place, a cost and an expectation of the first visit	The district programme's information component, and any helpline named in the message	Behaviour rather than report, ideally service contact attributable to the campaign
First response	What to do, and what not to do, before care is reached	Trained workforces and community groups around the person at risk	Knowledge in the responder, and where possible what the responder then did

Public awareness is a named component of the district programme, so the mandate exists everywhere the programme runs; what is published nationally is activity rather than knowledge, and no Indian literacy outcome with a stated population and denominator is established.

Psychiatric rehabilitation and community care

21 · Principles and components of psychiatric rehabilitation

Rehabilitation is what a service does about the part of illness that treatment does not reach. Symptoms can remit while a person remains unemployed, unhoused, unmarried, out of education and out of every role that gave their life shape, and no amount of additional medication addresses that. Psychiatric rehabilitation is the organised attempt to restore or build those roles, and read at population scale it is not an add-on to treatment but a distinct arm of the service with its own goals, its own workforce, its own outcomes and its own place in the architecture of a health system.

This section takes rehabilitation cross-diagnostically and at the level of systems: what the goals are, how disability and participation are framed and measured, who does the work, how the components are assembled into a network, and what has to happen for any of it to exist in a district. Rehabilitation as it is selected and delivered for an individual patient with schizophrenia, which is the disorder-specific psychosocial account, is taught in the Schizophrenia: management,

antipsychotics and adverse effects notes and is not restated here. Supported employment, deinstitutionalisation, the community settings and the Indian models each have their own section in this chapter and are named rather than taught here.

21.1 Three ideas that organise the whole field

Almost everything in rehabilitation follows from three constructs, and the reason candidates find the field vague is that they learn its activities without learning these.

- **Recovery** in this literature is not synonymous with remission. It denotes a person living a life they regard as worth living, with or without continuing symptoms, and it is defined by the person rather than by the clinician. That shift in who defines the endpoint is the single most consequential idea in modern rehabilitation, because it makes the goal negotiable rather than prescribed.
- **Disability** is the gap between what a person can do and what their environment demands and permits. It is therefore not a property of the person alone, and it can be reduced by changing the environment as readily as by changing the person, which is what makes it a public-health object rather than only a clinical one.
- **Participation** is the outcome that both of the above are aiming at, and it has a usable definition: participation means joining in community activities. That definition is deliberately mundane, and its mundanity is the point. It is checkable, it is meaningful to the person and their family, and it does not require an instrument to notice.

The international orientation of the field follows the same logic. **The World Health Organization's guidance on community mental health services sets out the case for shifting away from institutional, biomedical models of care and toward person-centred, recovery-oriented and human-rights-based approaches**. Each of those three adjectives corresponds to one of the constructs above, and together they describe not a technique but an orientation that a whole service either has or does not.

RULE **The rehabilitation goal is a role, not a score.** A plan whose stated aim is symptom reduction or an improved rating is a treatment plan wearing a rehabilitation label. A rehabilitation aim names the role to be resumed or acquired, the setting in which it happens, and the supports required to sustain it, and it is agreed with the person whose role it is.

21.2 The disability frame, and why psychiatry borrows it

Psychiatry does not need a private theory of disability, and using the general health one has three advantages: it is cross-diagnostic, it is shared with the rest of medicine and with the disability sector, and it is the frame in which entitlements are written. **The International Classification of Functioning, Disability and Health supplies that framework, and it is explicitly designed to apply at both the individual and the population level**, which is precisely what a chapter about services needs: the same construct that describes one person's limitation aggregates into a description of a district's burden.

Its companion assessment instrument is the **World Health Organization Disability Assessment Schedule, WHODAS 2.0**, and its defining property for our purposes is scope: it is used across all diseases, including mental, neurological and addictive disorders. That cross-diagnostic reach is why a rehabilitation service can use one functioning framework for a mixed caseload rather than a different measure for every diagnosis, and it is why functioning data from a psychiatric service can be compared with, and pooled with, functioning data from the rest of the health system.

The instrument is named here and not reproduced. Its domain structure, its scoring and its administration belong to the chapter that owns assessment instruments, the Psychotherapies, clinical assessment, MSE and nosology notes, and no part of it is printed in this chapter. What a resident needs from it here is the conceptual point: functioning is measurable, it is measurable in the same terms across conditions, and it is a different quantity from symptom severity.

21.3 Community inclusion as the outcome, and as an entitlement

Set against a history in which people with mental illness were removed from society for years at a time, the modern statement of the rehabilitation aim is unusually direct. In India it is not merely a professional aspiration but a legal one: **under the Mental Healthcare Act, 2017, every person with mental illness has the right to live in, be part of, and not be segregated from society**. The architecture of that statute, the machinery of rights, capacity, admission and review that surrounds the provision, belongs to the Mental health legislation and forensic psychiatry notes and is not rehearsed here.

What matters for a service is the direction of the obligation this creates. If community living is an entitlement, then segregation requires justification rather than inclusion requiring permission, and a person occupying an institutional bed for want of anywhere to go is not a clinical fact but a service failure with a legal character. That reframing is what converts rehabilitation from a benevolence into a duty, and it is the argument a resident should be able to make in one sentence.

IN INDIA

Two Indian facts sit together and they generate a concrete decision. Community living is a statutory entitlement, and the national programme's stated objectives include promoting community participation in the development of mental health services and stimulating efforts toward self-help in the community. Together they mean that when a resident writes a rehabilitation plan in a public service, the community component is not discretionary and is not something to be added if resources allow. The operational consequence is that the plan should name the community role, the person or group in the community who will support it, and the district programme staff member responsible, before it names any additional treatment.

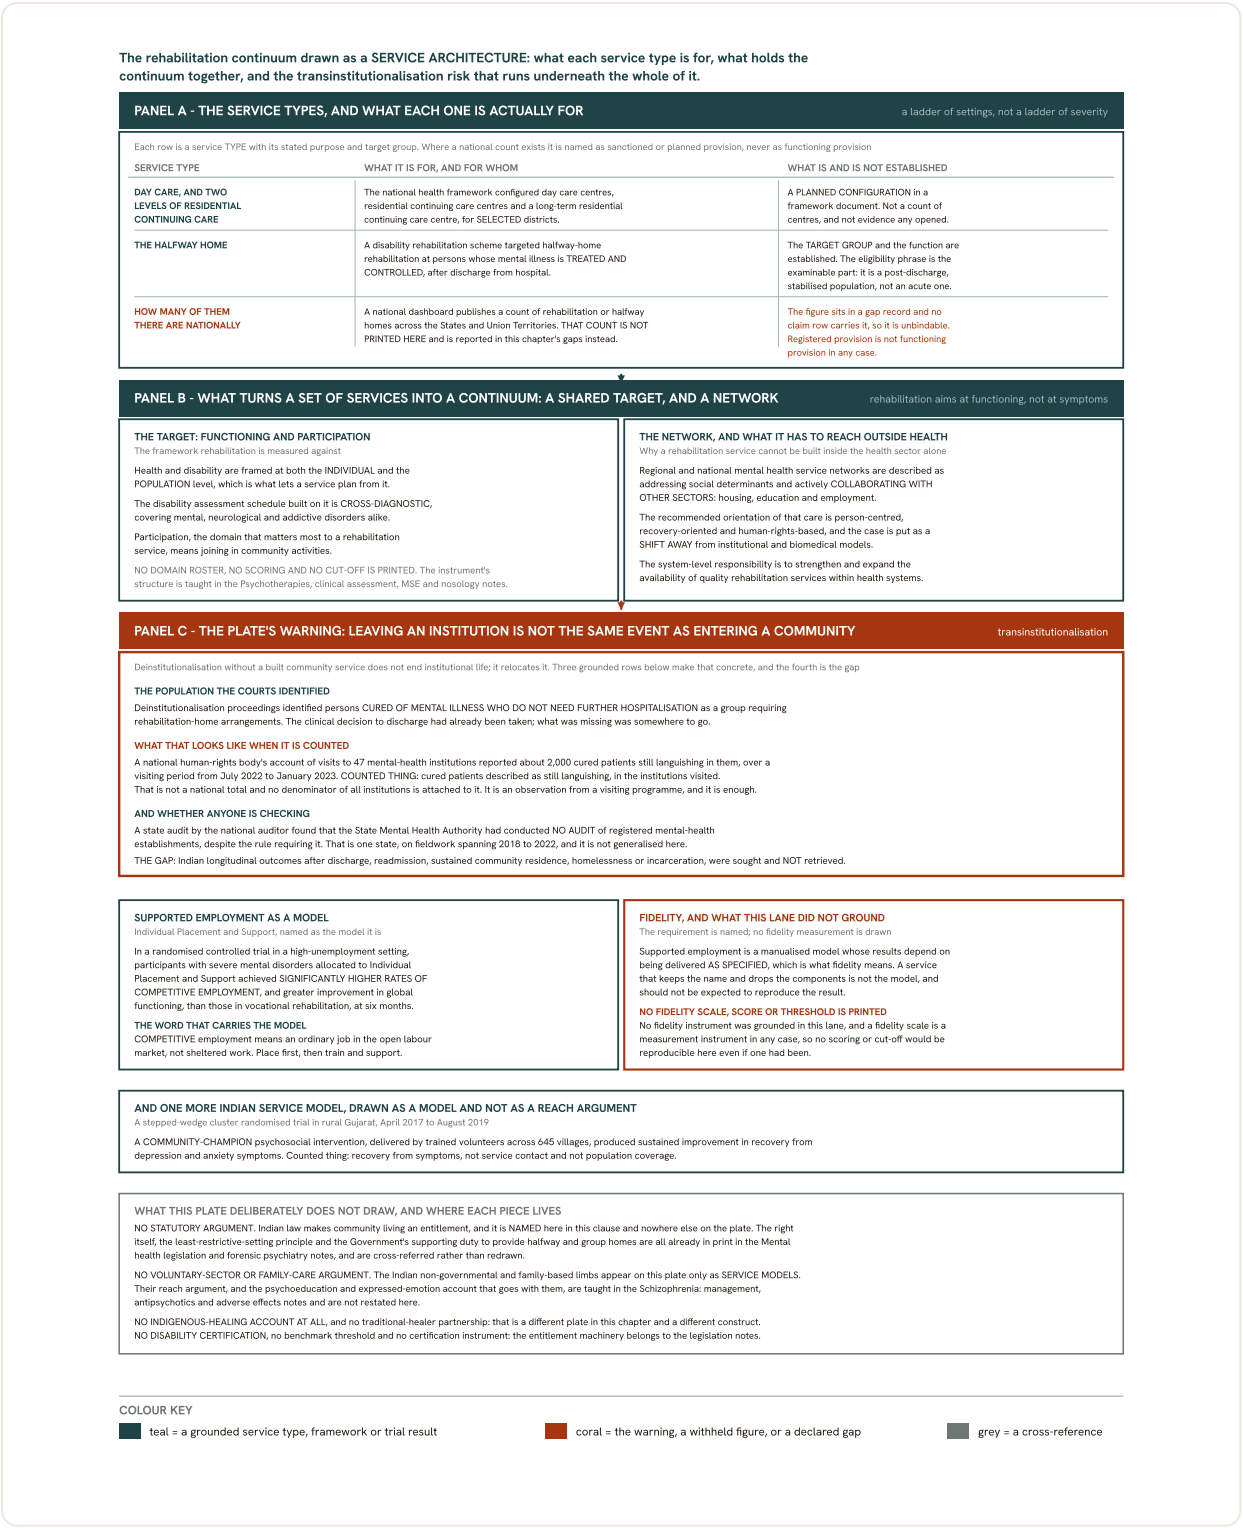


Fig 39.8 The rehabilitation continuum as a service architecture, with transinstitutionalisation as its warning. The national framework configured day care and two levels of residential continuing care for selected districts, and a disability scheme targeted halfway-home rehabilitation at people whose mental illness is treated and controlled after discharge; the national count of such homes sits in a gap record that no claim row carries, so it is not printed. The continuum is held together by a target, participation, meaning joining in community activities, and by service networks that collaborate across housing, education and employment. The warning is concrete: courts identified cured people who no longer needed hospitalisation but had nowhere to go, a human-rights body reported about 2,000 such patients still in the 47 institutions it visited, one state auditor found no audit of registered establishments had been done at all, and Indian post-discharge outcomes were sought and not retrieved. Supported employment is drawn as a model: it achieved significantly higher competitive-employment rates than vocational rehabilitation in a randomised trial, and its fidelity requirement is named without any fidelity instrument being reproduced.

21.4 The components: three ways to close the same gap

If disability is the gap between what a person can do and what the environment demands, then only three kinds of action can close it, and every rehabilitative component in the field is one of the three. This is the most useful classification a resident can carry, because it converts a shapeless list of activities into a decision.

Raise the person's capacity. Skills training in the domains the role requires, whether that is self-care, household management, money, travel, conversation, or the specific competencies of a job or a course. This is the component people think of first, it is the slowest of the three, and it is the only one that transfers with the person when they move.

Supply a support that substitutes for the missing capacity. A person, a structure or a piece of assistance that does what the person currently cannot: a family member who manages medication, a supported setting, a job coach, a case manager, a prompt, an entitlement that replaces lost income. Supports work immediately and stop working when they are withdrawn, which is why the plan has to state whether a support is a bridge or a permanent adjustment.

Lower what the environment demands. Modified hours, an adjusted course load, a changed task, a physical relocation of a service closer to where people live, a form simplified, a rule waived. This is the cheapest of the three and the one clinicians reach for last, because it requires negotiating with somebody outside the health system rather than doing something to a patient.

Most real plans use all three, and the useful discipline is to say which is which, because they have different timescales and different failure modes. A plan built only of skills training will show nothing for months; a plan built only of supports will collapse when the support is withdrawn; and a plan built only of environmental adjustment leaves the person no more able than before if the environment changes again.

21.5 Who does rehabilitation, and what a team has to contain

Rehabilitation is the least doctor-centred part of psychiatry, and the composition of the team is not incidental to it. In India the district programme is built around a multidisciplinary team: **the district team under the programme includes a psychiatrist, a clinical psychologist and a psychiatric social worker**. Those three disciplines map onto three different problems, and understanding the mapping is more useful than memorising the list. The psychiatrist holds diagnosis, medication and risk; the clinical psychologist holds assessment of functioning and psychological intervention; the psychiatric social worker holds the person's relationship with their family, their entitlements and their community, which is where most rehabilitation actually happens.

Two riders follow. First, a team roster is a statement about establishment and not about occupancy: a sanctioned post is not a person, and a district with vacancies in two of the three disciplines has a team on paper only. Second, the roster above is what the programme names and is not asserted here as the complete list of anybody's staffing; occupational therapy, nursing, community health workers, peer workers and vocational staff all have roles in a mature rehabilitation service, and the absence of any of them shapes what that service can offer.

TRAP **Listing rehabilitation activities in place of rehabilitation principles.** Asked about psychiatric rehabilitation, most candidates produce a list of interventions. An examiner is looking for the organising logic first: a person-defined goal expressed as a role, disability understood as the fit between person and environment, participation as the outcome, and a team whose composition follows from the fact that most of the work is social. The activity list is worth marks only once that frame is on the page.

21.6 From a team to a system: networks and other sectors

A team can rehabilitate a person only as far as its own walls. Everything a person needs to occupy a role, somewhere to live, something to do, money, transport and a social world, is held by other sectors, which is why rehabilitation is the part of psychiatry most dependent on institutions that do not report to health. This is not a rhetorical point but the accepted design: comprehensive mental health service networks are described as addressing the key social determinants of health and actively collaborating with other sectors such as housing, education and employment.

The practical form of that collaboration is unglamorous and specific: a named contact in the social welfare office, a working relationship with a disability certification process, a school or college willing to take a student back on modified terms, an employer prepared to make an adjustment, a landlord or a hostel that will accept a tenant. A rehabilitation service is largely a set of such relationships, and the ones that are written down and institutionalised outlive the individuals who built them.

21.7 Getting it to exist: implementation as the real problem

Rehabilitation suffers from a characteristic implementation deficit. It is uncontroversial, it appears in every policy document, and it is the first thing cut when a service is under pressure, because its absence produces no immediate crisis. The international guidance addresses exactly this: the recommendations on rehabilitation in health systems are addressed to member states and other stakeholders in order to strengthen and expand the availability of quality rehabilitation services, which is to say the problem being solved is availability rather than technique.

What has to be true for a district rehabilitation service to exist, in order of how often each is missing: a named person whose job it is; a physical space that is not a ward; a link to at least one non-health sector; a referral route in and a discharge route out; and a record of who is in the service and what they are working toward. Notice that none of those is a treatment, and that a service with all five and no new intervention will do more rehabilitation than a service with a manual and none of them.

21.8 The frameworks compared, and what each is good for

The commonest confusion in this topic is between the constructs, the frameworks and the instruments, so it is worth setting them side by side with what each can and cannot deliver.

WHAT IT IS	WHAT IT FIXES	LEVEL IT OPERATES AT	WHAT IT CANNOT DO
Recovery orientation	Who defines the goal, and therefore what counts as success	The individual, and the culture of a whole service	Be measured directly by any symptom or disability score
The functioning and disability classification	A shared vocabulary for limitation and for the environment's contribution to it	Individual and population, by design	Tell you what to do about any particular limitation
The cross-diagnostic disability assessment schedule	Comparable measurement of functioning across conditions	The individual, aggregable upward	Substitute for a symptom measure, or for the person's own account of their goals
The rights framework in statute	The direction of the obligation: inclusion is the default and segregation needs justifying	The service and the state	Create the community capacity that makes the entitlement usable
Service networks with other sectors	Access to housing, education, employment and income, which sit outside health	The system	Function without named relationships and someone to maintain them

GOOD TO REMEMBER

In an outpatient clinic the whole of this section reduces to one extra question at the end of a review: what does this person do all day, and what would they like to be doing that they currently cannot? A person whose symptoms are controlled and whose answer to that question has not changed in two years is not doing well, whatever the mental state examination records, and that is the moment to write a rehabilitation aim rather than to adjust a dose.

21.9 Evaluating rehabilitation, and the honest position

Rehabilitation is evaluated on roles occupied and sustained rather than on symptoms: living situation, occupation or education, income and entitlements received, social contact, and functioning measured on a cross-diagnostic scale. Two design cautions apply. Outcomes take time, so a study window shorter than the time it takes to acquire a role will find nothing regardless of the intervention. And the environment is part of the outcome, so a service operating in a labour market or a housing market that offers nothing will produce poor role outcomes with excellent practice, which is an argument for reporting the setting alongside the result rather than for despair.

MNEMONIC

Goal, Gap, Team, Sector, Place

The five things a rehabilitation plan has to name. The person's own goal expressed as a role; the gap between what they can do and what the environment demands; the discipline in the team who will hold the work; the sector outside health whose cooperation is needed; and the physical place where it happens. A plan missing any one of them will not survive contact with a district.

Role	**Cross-diagnostic**	**Statutory**	**Availability**
THE UNIT OF A REHABILITATION GOAL, NOT A SYMPTOM SCORE	SCOPE OF THE DISABILITY ASSESSMENT FRAMEWORK USED	STATUS OF COMMUNITY LIVING FOR PEOPLE WITH MENTAL ILLNESS IN INDIA	THE PROBLEM THE INTERNATIONAL GUIDANCE ADDRESSES, NOT TECHNIQUE

Rehabilitation aims at a role the person chooses, works on the gap between the person and their environment, measures functioning on a framework shared with the rest of medicine, and in India rests on a statutory entitlement to live in and not be segregated from society.

22 · Individual Placement and Support (IPS) and supported employment

Work is the rehabilitation outcome patients ask for and services most often decline to organise. Ask a person with a severe mental illness what they want and paid work appears near the top of the list far more often than clinicians expect. Work supplies income, structure, identity, social contact and a reason to be somewhere at a fixed time, all of which are exactly the things illness takes away, and it is the outcome most visible to the family. Supported employment is the service class built to deliver it, and Individual Placement and Support is the model that dominates the field.

This section treats the model as a programme: how a vocational service is designed, what its implementation barriers are, what its fidelity means, how it is evaluated, and what its Indian position actually is. The model as an intervention chosen and delivered for one patient with schizophrenia, and the older distinction between preparing a person before placing them and placing them and then supporting them, are taught in the Schizophrenia: management, antipsychotics and adverse effects notes and are named here rather than re-derived. The principle set and fidelity criteria published in the model's own manual are not reproduced in this chapter.

22.1 Why a health service is in the employment business at all

The objection is worth answering directly, because an examiner may raise it. Employment is not a health service's core competence, employers are not health institutions, and a labour market is not something a district hospital can influence. Three arguments answer it.

- **The barrier is clinical in origin.** The reason many people with severe mental illness are out of work is not a lack of jobs in general but a set of consequences of illness and its treatment,

together with discrimination. Nobody else is positioned to address those, and a generic employment service that does not understand them will fail with this population.

- **The support has to be coordinated with treatment.** Medication timing, appointment scheduling, relapse recognition and the handling of a period of illness while employed all require the clinical team, and a vocational service operating with no line to the clinicians will lose people at the first setback.
- **Work changes the clinical picture.** Income, structure and role are not incidental to outcome, which is why employment is a rehabilitation aim and not a life event that happens after treatment finishes.

RULE **The endpoint is competitive employment, not activity.** **Competitive employment** means an ordinary job in the ordinary labour market, open to anyone, paid at the going rate. Occupational activity within a service, sheltered production and unpaid work experience are different endpoints entirely, and a programme that reports them as employment has changed the outcome definition without saying so.

22.2 The two service logics, and what each demands of a system

Vocational rehabilitation has historically been organised around preparing a person first, through assessment, training and graded work in a protected setting, and placing them once they are judged ready. **Supported employment** inverts the order: the person is helped into an ordinary job early and the support is delivered on the job. The contrast itself is taught in the notes named above; what belongs here is what each logic requires a service to build, because that is what makes one of them feasible in a district and the other not.

A preparation-first system requires physical infrastructure, a sheltered setting, equipment, supervision and a production process, and its unit of investment is a building. A place-and-support system requires almost no infrastructure and a great deal of relationship: employers who will take a referral, a specialist who can spend the day travelling, and a clinical team willing to be interrupted. The first is easier to fund because it is visible and countable; the second is harder to fund and reaches ordinary employment.

The distributional consequence is the one worth stating in an answer. A preparation-first system has an implicit gate, since somebody decides who is ready, and gates exclude the people with the most severe illness. A place-and-support system with an open door reaches those people and accepts that some placements will fail. Which of those a service prefers is a value choice about who the service is for, not a technical one.

22.3 The design decisions that define a supported-employment service

Individual Placement and Support is the best-specified expression of the place-and-support logic. Rather than reciting its published principles, it is more useful to set out the decisions any vocational service has to make, because those decisions are what a resident will be asked about and what a service actually gets wrong.

DESIGN DECISION	THE RESTRICTIVE OPTION	THE SUPPORTED-EMPLOYMENT OPTION	WHAT THE CHOICE COSTS OR BUYS
Who is eligible	Referral after an assessment of readiness	Anyone who says they want to work	An open door reaches severe illness and raises the failure rate; a gate protects the success rate by excluding the people the service exists for
What counts as a job	Sheltered or in-service occupational activity	An ordinary paid job in the open labour market	Sheltered work is easier to arrange and does not produce the income, status or social contact that the outcome is valued for
When the search starts	After a period of preparation and training	Early, with support delivered on the job	Preparation delays and loses people at every step; early search accepts that learning happens in the real setting
Whose preference decides	What the service can arrange	The kind of work the person actually wants	Jobs chosen by the person are held longer; jobs chosen for them are convenient to arrange and quickly lost
Where the employment worker sits	In a separate vocational agency	Inside the clinical team, sharing its meetings and its record	Integration is the expensive part and is repeatedly the part dropped, after which the service is no longer the model it is named for
How long support lasts	Until placement, then discharge	Continuing as long as it is needed, on both sides of the placement	Time-limited support produces placements and not retention; continuing support is a standing cost with no dramatic moment to justify it
Whether income advice is included	Not the service's business	Part of the work, because entitlements interact with earnings	Without it a person can be advised into a job that leaves them worse off, and will reasonably refuse the next offer

Reading the table downward gives the model. Reading it across gives the examination answer, which is that each decision has a defensible restrictive version and that the supported-employment position is a coherent set of choices rather than a slogan.

22.4 The evidence, and how to read a single trial

The comparison that matters is against conventional vocational rehabilitation rather than against nothing. In a randomised controlled trial conducted in a high-unemployment setting, **participants receiving Individual Placement and Support achieved significantly higher rates of competitive employment and showed greater improvements in global functioning than participants receiving vocational rehabilitation**, with those outcomes assessed at six months in people with severe mental disorders who met the trial's eligibility criteria.

Four things in that sentence are load-bearing and a candidate should be able to point to each. The comparator is an active one, so the finding is not merely that doing something beats doing nothing. The primary outcome is competitive employment, which is the endpoint that matters rather than a proxy. A second outcome, global functioning, moved in the same direction, which is the finding that argues against the objection that employment gains come at a cost to the person's mental state. And the setting was one of high unemployment, which matters for the next subtopic.

TRAP **Quoting an employment rate as though it were a property of the model.** Employment rates achieved in vocational trials are functions of the model, of the eligibility criteria and of the labour market the trial ran in, and the last of the three moves the number more than the first. A rate quoted without its labour market is uninterpretable, which is why this section reports the direction of the comparison and prints no employment percentage at all.

22.5 Transportability, which is the real Indian question

Transportability, meaning whether a result obtained in one place will hold in another, is the question this evidence raises most sharply, because every employment outcome is defined against a labour market and the first thing to ask of any vocational evidence is what market it came from. A trial in a high-unemployment setting is informative in a specific way: it tests the model where jobs are scarce, which is a harder test than a tight labour market provides, and it therefore travels better than a result obtained where employers are competing for staff.

The Indian question is not only about the quantity of work but about its form. Where a large share of employment is informal, seasonal, self-organised or family-based, the definition of competitive employment, an ordinary job in the open market at the going rate, describes a smaller part of the working population than it does in the settings where the model was developed and tested. That does not make the model inapplicable; it means the outcome definition has to be examined before the model is imported, and that a service which counts only formal-sector placements will under-record what it achieved. Work within a household or a family enterprise is a real vocational outcome in this setting, and it is taught in the notes cross-referred above rather than here.

The honest position on Indian implementation is that this chapter's sources carry no coverage figure, no service count and no outcome data for supported employment in India. Formal vocational services organised on this model are scarce, and a resident asked about Indian supported employment should say that the model's evidence base is international, that no

national Indian implementation figure is established, and that the transportability question is about the labour market and the outcome definition rather than about the intervention.

22.6 Fidelity, and why the word matters

Fidelity is the degree to which a service as delivered matches the model as specified. It matters in this field more than in most, because supported employment is easy to name and hard to run, and the components that get dropped under pressure are systematically the expensive ones: the integration of the employment worker into the clinical team, the open door, and the continuation of support after placement. A service that has dropped those is still called by the model's name, and its results are then reported as results of the model.

Two consequences follow. For a reader of the literature, an implementation study that reports poor outcomes without reporting fidelity has not shown that the model failed. For a service, fidelity is the thing to measure when outcomes are disappointing, before concluding that the approach does not work locally. Structured fidelity instruments exist for this model and are published by the groups that developed them; none is reproduced here, and a service intending to use one should obtain it from its rights holder.

> PITFALL
>
> Calling a job placement scheme by the model's name because it places people in jobs. The name denotes a specified set of components, and the two that are almost always missing in practice are the employment worker's membership of the clinical team and the continuation of support after the person starts work. A service without those places people and then loses them, which is not the model performing badly but a different service performing as expected.

22.7 Evaluating a vocational service

The endpoints form an ordered set and each answers a different question. Number entering the service answers a reach question. Number placed in competitive employment answers an efficacy question. Time to first placement answers an efficiency question. **Job tenure** and total time worked answer the retention question, which is the one that distinguishes a service from an introduction agency. Earnings answer the question the person actually asked. And functioning and clinical outcome answer the objection that work will destabilise people.

Two cautions apply to all of them. The denominator has to be stated, since a placement rate computed over people who completed a preparation phase is not comparable with one computed over everybody who walked through the door. And the follow-up has to be long enough for retention to be visible, because a placement measured at the moment it happens tells you nothing about whether the person is still there.

There is one further evaluation trap peculiar to this field and it is worth naming. Employment outcomes are highly sensitive to who is counted and when the counting starts, so two services with identical practice can report very different results by defining their cohort differently: everyone referred, everyone who attended once, everyone who set a work goal, or everyone still engaged at the point of analysis. The last of these is the most flattering and the most common,

because it silently removes the people for whom the service did not work. A published vocational result should always be read backwards from the denominator to the referral, and a service reporting its own results should state the referral cohort before anything else.

GOOD TO REMEMBER

Ask about work at every review, in the same routine way as sleep and appetite, and ask what the person wants rather than what they are ready for. In a service with no vocational specialist the useful actions are small and available: agree a work goal in the notes, schedule appointments so they do not collide with a working day, avoid regimens that make morning work impossible where an alternative exists, and speak to the family about the difference between protecting somebody and keeping them idle.

MNEMONIC

Open door, Ordinary job, Own choice, One team, Ongoing support

Five things a supported-employment service either has or has not. No readiness gate, an ordinary paid job in the open market, the work the person chose, the employment worker inside the clinical team, and support that does not stop at placement. When outcomes disappoint, check these before questioning the approach.

Competitive THE ONLY EMPLOYMENT ENDPOINT THAT COUNTS AS THE OUTCOME	**Higher** EMPLOYMENT WITH THE MODEL AGAINST VOCATIONAL REHABILITATION, RANDOMISED COMPARISON	**Retention** WHAT SEPARATES A SERVICE FROM AN INTRODUCTION AGENCY	**Not established** ANY INDIAN COVERAGE OR OUTCOME FIGURE FOR THIS MODEL

Supported employment places the person in an ordinary paid job first and supports them there, its endpoint is competitive employment and its evidence is a randomised comparison against conventional vocational rehabilitation, and what fails in practice is almost always fidelity rather than the model.

23 · Deinstitutionalisation: rationale, outcomes and lessons

A service transition, not an event and not a policy announcement. **Deinstitutionalisation** names the deliberate shift of the centre of gravity of psychiatric care out of large custodial institutions and into general hospitals, community teams and ordinary housing. Every country that has attempted it has discovered the same thing: closing beds is easy, and building the community capacity that has to receive the people is slow, expensive and politically invisible. The distance between those two facts is where the entire teaching content of this topic lives.

The statutory right that mandates community living, and the rights architecture that surrounds it, belong to the Mental health legislation and forensic psychiatry notes and are named here rather than rehearsed. What this section owns is the transition itself: why it was undertaken, what its outcomes were, what has to exist for it to succeed, what happens when the second half is not

built, and what the Indian record actually documents. The community settings that receive people, the day care centre and the halfway home, are configured in their own section of this chapter.

23.1 What drove it, and why the drivers matter

Four forces converged, and it is worth separating them because they pull in different directions and because a service that understands only one of them will make characteristic mistakes.

- **Therapeutic.** Effective treatment shortened the period during which continuous institutional supervision was necessary, so a category of person appeared who no longer required a hospital and had nowhere else organised to be.
- **Ethical and legal.** Long-stay institutions were repeatedly found to produce harm of their own, and the growth of rights-based thinking made indefinite custodial care without justification indefensible rather than merely regrettable.
- **Clinical.** Prolonged institutional living produces its own disabilities, a loss of initiative, of skills and of social contact, which are additional to the illness and are caused by the setting.
- **Economic.** Large institutions are expensive to run, and this is the driver that most often becomes the operative one. That matters, because a transition motivated by cost has a strong incentive to complete the closure and a weak incentive to fund the replacement.

RULE **Deinstitutionalisation is a two-sided programme and only one side is self-executing.** Bed closure happens by administrative decision. Community capacity happens only if somebody funds, staffs and builds it. A programme that treats the second as a consequence of the first has not planned a transition, it has planned a discharge.

23.2 The arc, told as a pattern rather than as a chronology

The international experience repeats a recognisable sequence, and knowing the sequence is more useful than knowing any one country's dates. It begins with a long asylum era in which the institution is the service and admission is effectively indefinite. A critique then accumulates from several directions at once: from clinicians who observe that a portion of the disability they are treating is produced by the setting, from lawyers and rights advocates who object to indefinite detention without review, from the emergence of treatments that make continuous supervision unnecessary for many people, and from finance ministries that notice the running cost of large estates.

Bed reduction then proceeds, usually faster than the community programme that was supposed to accompany it, and the visible consequences appear in other sectors rather than in psychiatry: in homelessness, in criminal justice, in nursing and residential homes, and in the burden carried by families. A later corrective phase follows in which community teams, crisis services, supported housing and assertive follow-up are built, and the systems that build them find that community care is not obviously cheaper than the institutions it replaced, only better, which is a very different argument to make to a treasury.

The clinical construct that justified the whole enterprise deserves its own name. **Institutionalism** is the syndrome produced by prolonged institutional living itself, comprising loss of initiative and of ordinary decision-making, submissiveness to routine, loss of practical and social skills through disuse, and a narrowing of expectation until the person no longer wants what they cannot have. It is additional to the illness, it is caused by the setting, and it is partly reversible, which is why it is an argument for moving people rather than an argument that they cannot be moved.

23.3 What has to exist on the receiving side

The precondition list is short and each item is a place where real programmes have failed.

Somewhere to live. Housing is the precondition, not an adjunct, and it is the one a health department cannot supply on its own. A discharge plan without an address is a transfer of the person to their family or to nowhere.

An income. A person leaving a long stay has usually lost earning capacity, savings and often any claim on family property. Without an income or an entitlement, community living is not sustainable regardless of clinical state.

A team that follows. Continuity of clinical care after discharge, with somebody responsible for finding the person when they do not attend, is what distinguishes community care from discharge. The absence of an **assertive follow-up** function, meaning a service that goes after the person rather than waiting for them, is the commonest technical failure.

Something to do. Occupation and role, addressed elsewhere in this chapter, without which the community placement is a smaller and lonelier institution.

A family or a substitute for one. In a system where families provide most long-term care, the family's willingness and capacity is a service variable and not a private matter, and it has to be assessed rather than assumed.

23.4 Transinstitutionalisation and the other failure modes

Transinstitutionalisation is the term for what happens when the receiving side is not built: people do not enter community life, they move into another institution. The destinations differ by country and by era, and the categories to know are prisons, nursing and old-age homes, custodial or beggar homes, and unregulated private facilities, with homelessness as the outcome when no institution absorbs them at all.

The concept is examinable because it explains why bed-count data can look like success while the population's situation has not improved. A country that halves its psychiatric beds and doubles the psychiatric population of its prisons has not deinstitutionalised anybody; it has changed which ministry is responsible. The corollary for evaluation is that the outcome measure cannot be a psychiatric bed count, and has to be a description of where people actually are.

TRAP **Treating bed reduction as the outcome of deinstitutionalisation.** Bed numbers measure the supply side and are easy to obtain, which is why they get reported. The outcome of interest is where the people are living, whether they are still in contact with services, and whether they have an income and a role. If those cannot be reported, the honest statement is that the programme's inputs are known and its outcomes are not.

23.5 What the Indian record actually documents

India's deinstitutionalisation story is not primarily a story of planned bed closure. It is a story about people who no longer need to be in hospital and are still there, and it has been documented mainly through litigation and inspection rather than through programme evaluation. Three kinds of record exist and each is a different sort of evidence.

The first is judicial. Deinstitutionalisation proceedings before the Supreme Court identified, as a population requiring rehabilitation-home arrangements, persons living with mental illness who have been cured and who do not need further hospitalisation. Note the construction of that category. It is defined administratively, by two negatives, no longer ill in the relevant sense and no longer needing a hospital, and it says nothing about what the person can do or wants. A category defined that way identifies who should not be where they are without identifying where they should go, which is precisely the gap the proceedings were addressing.

The second is inspectorial. **An account by the National Human Rights Commission of visits to 47 mental-health institutions reported about 2,000 cured patients still languishing in institutions**. Read the scope carefully before using the figure. It describes patients in the institutions that were visited, so it is not a national census, it carries no denominator, and no proportion may be computed from it. What it establishes is that the phenomenon exists at institutional scale and has been found by an official body that went and looked.

The third is audit. A performance audit found that a state mental health authority had not conducted the audit of registered mental-health establishments that it was required to conduct. That finding is about the machinery rather than about patients, and it is the most quietly damaging of the three, for the reason set out below.

47	about 2,000	None	Not established
MENTAL-HEALTH INSTITUTIONS VISITED IN THE INSPECTION ACCOUNT	CURED PATIENTS STILL IN THOSE INSTITUTIONS, NOT A NATIONAL CENSUS	VALID DENOMINATOR FOR THAT COUNT, SO NO PROPORTION	WHAT HAPPENS TO PEOPLE AFTER DISCHARGE IN INDIA

23.6 The accountability machinery, and why the audit finding matters

A system that intends to move people out of institutions needs to know what its institutions are doing, which is why regulatory inspection of registered establishments exists as a duty rather than as a good idea. When an audit finds that the required inspection of registered mental-health establishments has not been carried out, the consequence is not merely a procedural lapse. It means that the state's own information about conditions, about length of stay and about who is ready for discharge is not being generated, and a transition cannot be planned on information that is not collected.

This is the general form of a lesson that recurs across this chapter: the failure is usually not at the level of policy, which is generally sound, nor at the level of clinical practice, which is generally willing, but at the level of the machinery that is supposed to observe and report. A resident who can identify the missing observational step in any programme is answering these questions at the right level.

IN INDIA

The operational decision is what to do when you have a patient who is ready for discharge and has nowhere to go. Three things follow from the record above and each is actionable. Record readiness for discharge explicitly in the notes on the day it is reached, because the category of person who no longer needs hospital care is the category the courts and the inspectorate act on, and an undocumented readiness is invisible to both. Involve the social work function early, because the obstacle is an address, an income and a receiving family rather than a treatment. And ask whether your establishment has been inspected by the state authority, because where that machinery is not operating, nothing about your ward's long-stay population is being reported to anybody.

23.7 What India does not know about its own discharges

Here the chapter has to be plain. No official Indian longitudinal or comparative data on what happens to people after discharge from mental-health institutions was retrieved. That includes readmission, sustained community residence, homelessness, incarceration, transfer into custodial or beggar homes, and every other transinstitutionalisation outcome. It was sought from the courts, the health ministry, the human rights commission, the national auditor, the disability empowerment department and state governments.

That absence is not a small one, and it is worth stating what it means rather than merely noting it. Without post-discharge outcome data, nobody can say whether Indian deinstitutionalisation, to the extent that it has happened, has produced community living or transinstitutionalisation. The country has documentation of the problem, judicial direction about it, and no measurement of the result. In a viva, a candidate who says exactly that, and who names transinstitutionalisation as the specific thing that would be measured, has given the strongest available answer.

PITFALL

Assuming that a discharge to family is the same as community living. In a family-based care system the family is the default destination, and for many people it is a good one. But a discharge that transfers the whole burden of supervision, income, occupation and crisis management to a household with no support, no training and no line to the service is a private institution with fewer safeguards than the public one. The question to ask before such a discharge is what the family is being asked to do, and what the service is offering them in return.

23.8 The lessons, stated as design rules

LESSON	WHAT GOES WRONG WHEN IT IS IGNORED	WHAT IT REQUIRES IN PRACTICE
Build the receiving capacity before reducing the sending capacity	People leave with nowhere to go, or do not leave at all and occupy beds that are no longer clinically needed	Housing, income and follow-up commissioned as part of the same programme, not as a later phase
Measure where people are, not how many beds there are	Bed reduction is reported as success while the population moves into other institutions	Post-discharge follow-up with named outcomes, including incarceration and homelessness
Treat readiness for discharge as a documented status	People ready to leave are invisible to the courts, the inspectorate and the planners	Explicit recording, and a register that somebody outside the institution can read
Keep the inspection machinery working	The state stops knowing what its own establishments contain	The statutory authority actually conducting the audits it is required to conduct
Fund the family as a service partner	Discharge becomes an unsupported transfer of custodial responsibility to a household	Carer support, crisis access and a named contact for the family

MNEMONIC

Roof, Rupees, Role, Review

The four things a person needs on the day they leave a long stay. Somewhere to live, an income or entitlement, something to do, and a clinical contact who will notice if they disappear. A discharge plan missing any one of the four is the mechanism by which deinstitutionalisation becomes transinstitutionalisation.

Deinstitutionalisation succeeds or fails on the receiving side, its true outcome measure is where people are living rather than how many beds were closed, and in India the long-stay population has been documented by courts and inspectors while what happens after discharge has never been measured at all.

24 · Community mental health services, day care and halfway homes

Between the ward and the family home there is a ladder, and most Indian districts own only its two ends. A person recovering from a severe illness rarely moves in one step from inpatient care to independent living, and the settings that occupy the middle of that ladder, day care, the halfway home, supported accommodation, longer-term residential care and the community team that visits, exist to make the intermediate steps possible. This section is about how each of those settings is configured: what it is for, who enters it, who staffs it, what a person does there, how they leave, and which arm of government pays for it.

The settings themselves will already be familiar. Their compact definitions, and the place each occupies on the rehabilitation ladder, are taught in the Schizophrenia: management,

antipsychotics and adverse effects notes, and the statutory right to community living and the prohibition on keeping a person in an institution who does not need to be there belong to the Mental health legislation and forensic psychiatry notes. Neither is repeated here. What follows takes the settings as known and asks the service questions about them, which is the part no other chapter carries.

24.1 What it means to configure a setting rather than to define it

A setting is defined by what it is called and configured by five decisions, and it is the five decisions that determine whether it works.

- **Function.** What is the setting actually for, in one sentence: containment, structure, skill acquisition, supervision of treatment, respite for a family, or transition to somewhere else.
- **Entry criteria.** Who is eligible, who decides, and what a person must be able to do before they are accepted.
- **Staffing.** Which disciplines, at what density, present at which hours.
- **Duration and exit.** How long a person stays, what triggers the move on, and where they move to.
- **Funding and governance.** Which arm of government or which body pays, registers and inspects it.

The fourth is the one that separates the settings from each other conceptually, and the fifth is the one that separates them administratively in India, which is why the last two subtopics of this section are about money and counting rather than about clinical content.

RULE **A community setting without a defined exit becomes a small institution.** Every intermediate setting is justified by the step it enables. Where no exit criterion is written and no destination exists, length of stay drifts upward, the population ages in place, and the setting acquires exactly the long-stay character the ladder was built to avoid. Write the exit before the admission.

24.2 The community team, which is the setting without a building

Before any of the physical settings, the ladder needs the service that connects them, and it is the one most often left out of an answer because it has no address. A community mental health team is defined not by where it sits but by where it goes: it holds a caseload of people living at home or in supported settings, sees them where they are rather than where the hospital is, and takes responsibility for noticing when somebody stops attending.

Its four functions are worth naming separately. It supplies *continuity*, so that a person is followed by a service rather than by a series of unconnected appointments. It supplies *outreach*, meaning contact that is initiated by the service, which is the single feature that distinguishes community care from an outpatient clinic and the one that reaches people who default. It supplies *coordination*, holding the relationships with the day-care centre, the residential setting, the family, the employer and the welfare office. And it supplies *crisis access*, without which every deterioration becomes an emergency admission.

In India this function is carried by the district mental health programme's team, and it is worth being clear about the implication. Where that team's posts are vacant or its outreach component is not funded, the entire ladder above and below it becomes a set of disconnected facilities, because there is nobody whose job it is to move a person along it. The physical settings are visible and get built; the connecting function is invisible and gets cut, and its absence is the commonest reason a district with facilities still has no community service.

24.3 Day care, configured

The **day care** centre is the least expensive intermediate setting because it borrows the person's own accommodation. Its function is structure and occupation during the day, together with a regular point of clinical contact, for a person who lives at home. Its natural population is a person whose family can hold the nights and the weekends but cannot supply a day that has any shape to it, which describes a very large share of the severe-illness population in a family-based system.

Its configuration questions are specific. Attendance is usually flexible, so a register that records who is expected and who came is the difference between a service and a room. Transport is the commonest reason for non-attendance and is almost never budgeted. Staffing needs to include somebody who can run an activity programme, not only somebody who can supervise medication. And because the person goes home each evening, the family is a partner in the day care arrangement rather than a bystander, which means the setting can be used deliberately to give a carer predictable hours of relief.

24.4 The halfway home, configured

The **halfway home** is residential and transitional, and both of those words are load-bearing. Residential, because it supplies accommodation and therefore reaches the person for whom home is not currently available or not currently workable. Transitional, because its whole justification is that the person is going somewhere afterwards.

Indian scheme design states the target group explicitly. **The Deendayal Disabled Rehabilitation Scheme of 2009 targeted halfway-home rehabilitation at people whose mental illness is treated and controlled, after their discharge**. Read the eligibility carefully, because it does two things at once. It defines the setting as post-discharge and therefore downstream of treatment, and it defines the population by clinical stability, which means the halfway home is not a place to send an unstable person and not a substitute for admission. A service that uses it as either has misconfigured it.

The configuration questions that decide whether such a home works are the length of stay and the destination. A halfway home with no ceiling on stay and no supported accommodation to move into becomes long-term residential care under another name, which is not a failure of the staff but a failure of the ladder above them.

TRAP **Describing the halfway home as a place for people who are unwell.** The scheme design defines its population as people whose illness is treated and controlled, after discharge. Candidates routinely describe it as an intermediate treatment setting for partially recovered patients, which imports an acute function it was not configured for and confuses it with day care and with continuing residential care in the same answer.

24.5 The continuing-care tier, and how India planned it

Not everybody moves on. A proportion of people with the most severe and persistent illness, particularly those with no family and no income, need residential support that is not transitional, and a ladder that pretends otherwise simply discharges them into homelessness. Indian programme design has recognised this in the structure of its own planning: **the National Health Mission implementation framework for 2012 to 2017 provided for day-care centres, residential continuing-care centres and a long-term residential continuing-care centre in selected districts**.

Three features of that plan are worth teaching. It is tiered, distinguishing day attendance from residential continuing care and again from long-term residential continuing care, which is exactly the ladder logic. It is explicitly limited to *selected* districts, so it was never a universal offer and should not be described as one. And it is a statement of what will be provided, which is a plan rather than an inventory, so nothing about how many such centres came to exist follows from it.

24.6 Two doors, two ministries, one patient

The single most useful administrative fact in this topic is that India's intermediate settings are reachable through two different arms of government, and they have different eligibility logics, different registering authorities and different money. The health route runs through the national and district mental health programmes and treats the setting as a health service. The disability and social welfare route runs through disability schemes and treats the same setting as a rehabilitation facility for a person with a disability, which brings certification, entitlements and a different inspection regime with it.

ROUTE	HOW THE SETTING IS FRAMED	TYPICAL ELIGIBILITY LOGIC	WHAT IT BRINGS WITH IT
Health programme route	A service component of the mental health programme, planned for selected districts as day care and as two tiers of residential continuing care	Clinical need, decided by the treating service	Clinical staffing, a link to the district team, and dependence on programme funding cycles
Disability and social welfare route	A rehabilitation project, with the halfway home as a named model project for people whose illness is treated and controlled after discharge	Post-discharge stability, and often a disability certificate	Grant-based funding to voluntary organisations, certification requirements, and a different inspectorate

The practical consequence for a resident is that the answer to "where can I send this person" depends on which door you knock on, and that the two doors have different waiting times, different paperwork and different reasons for refusing. Knowing that there are two is worth more in a district job than knowing the name of any single facility.

IN INDIA

The decision this section forces is procedural and it is made early or not at all. If a person is likely to need an intermediate residential setting after discharge, start the disability certification and the scheme paperwork while they are still an inpatient, because the social welfare route requires documentation that takes time to assemble and cannot be produced on the day of discharge. Simultaneously ask the district programme what it actually operates locally, since the health route's day-care and continuing-care provision was planned for selected districts rather than for all of them. Doing both in parallel is the only way to have a destination ready when the person is ready.

24.7 What is counted, and what a count of it means

A national dashboard publishes a count of rehabilitation homes and halfway homes across states and union territories, which sounds like the coverage figure this section needs and is not. This chapter cannot print those figures, because no claim row here asserts them, and the honest treatment is to teach what such a count would and would not establish.

A dashboard count is a count of *listed or registered* facilities. It is not a count of facilities that are open, not a count of places available, not a count of people accommodated, and it has no denominator, so no proportion of need met can be computed from it. Two facilities appearing as one row each may differ by an order of magnitude in the number of people they hold. And a national total says nothing about distribution, which is the variable that matters to a patient, because a home in another state is not a destination.

The deeper gap sits underneath the count. Current operative staffing standards and referral criteria for day-care centres and halfway homes, together with any national or state outcome data that would distinguish sanctioned or registered provision from functioning services, were sought across the health ministry and mission, the disability empowerment department and state health and social welfare authorities, and were not retrieved. So the position to state in a viva is this: the settings are named in Indian scheme design, their target populations are defined, a national listing of homes exists, and what they are staffed with, who they must accept and what happens to the people in them is not established.

PITFALL

Quoting a registry total as service capacity. It is the same error as reading a sanctioned district total as coverage, and it recurs across this chapter because registries are the only national data that exist for community services. A listed home may be closed, may be full, may take only one sex, may charge fees, and may be a day's travel away. Before referring, confirm the four things a registry never records: is it open, is there a place, what does it cost, and will it take this person.

Supported accommodation deserves a mention as the rung above the halfway home, because it is the destination that makes a transitional setting transitional. It is ordinary housing with a support component attached, ranging from a warden or visiting worker to a shared flat with staff on call, and its defining feature is that the tenancy rather than the treatment is the organising fact. Where it does not exist, everything below it silts up. This chapter's sources carry no figure for its availability in India, so the honest statement is that the rung is named in policy and its extent is unmeasured.

24.8 Choosing a setting for a particular person

The decision reduces to three questions asked in order. Does the person have somewhere to sleep tonight that is safe and available? If yes, the residential settings are not indicated and the question is what fills the day. If no, is the need transitional or continuing? A transitional need with a plausible destination points to a halfway home; a continuing need with no destination points to the continuing-care tier, and pretending otherwise produces a failed placement. And finally, who is going to keep clinical contact wherever the person goes? A community setting without a named clinical contact is where people are lost to follow-up.

MNEMONIC

Function, Entry, Staffing, Exit, Payer

The five configuration questions, in the order in which they should be asked of any community setting. If you can answer all five for the day-care centre and the halfway home in your own district, you know more about local services than most of the people referring into them.

Selected DISTRICTS THE PLANNED DAY-CARE AND CONTINUING-CARE TIERS WERE FOR	**Post-discharge** TARGET POPULATION FOR THE HALFWAY-HOME MODEL PROJECT	**Two** GOVERNMENT ROUTES INTO THE SAME SETTINGS, HEALTH AND SOCIAL WELFARE	**Not established** STAFFING STANDARDS, REFERRAL CRITERIA AND OUTCOMES FOR EITHER SETTING

Day care borrows the person's own accommodation and supplies the day; the halfway home is residential, transitional and configured for people whose illness is treated and controlled after discharge; and in India both are reachable through a health route and a social welfare route with different eligibility, different money and different inspectors.

25 · Models of care in India: NGOs, family-based care and indigenous resources

Most mental health care in India is delivered by people the health system does not employ. Families, voluntary organisations, community volunteers and traditional practitioners between them carry a far larger share of the load than the district programme does, and a chapter on service architecture that describes only the government machinery has described the smaller part of the system. This section treats those three as *models of care*: configured arrangements with an identifiable provider, a funding basis, a governance structure and, in principle, an evaluable output.

Two boundaries apply before we start. The family as the delivery vehicle for psychoeducation and for work with high expressed emotion, and the voluntary sector named as a force multiplier for community reach in the management of schizophrenia, are both taught in the Schizophrenia: management, antipsychotics and adverse effects notes, and neither is restated here. The design of a formal partnership with traditional and faith-based healers, meaning bidirectional referral, training, safeguards and access, is taught in the section of this chapter that owns traditional healers, and is named here rather than built.

25.1 What makes something a model rather than an activity

A **model of care** is not the same as a thing that helps. Four properties turn an arrangement into a model, and they are the properties that let it be commissioned, replicated and evaluated.

- **A defined provider and a defined role.** Somebody specific does something specific for somebody specific, and it is written down.
- **A funding and governance basis.** Money comes from somewhere on a stated basis, and somebody is accountable for how it is used.
- **A relationship with the formal system.** Referral in, referral out, and an agreed division of what each side does, without which the model is a parallel service rather than a component of one.
- **An output that could be measured.** Not necessarily measured, but of a kind that could be.

Running the three Indian limbs through those four properties is the most useful thing this section does, because it explains immediately why the voluntary sector is the best documented, why family care is the largest and least visible, and why indigenous resources are the hardest of the three to hold to any of it.

RULE **Name the provider, the money and the referral relationship, or you have not described a model.** A great deal of writing about Indian community psychiatry describes admirable activity without any of the three, which is why so little of it can be commissioned by anybody or repeated anywhere else.

25.2 The voluntary sector as a configured model

The **voluntary sector**, meaning non-governmental organisations working outside both the state system and private practice, occupies a specific and predictable position in Indian mental health care. They tend to do the things the public system finds hardest to fund: long-term residential rehabilitation, day activity, vocational work, homeless outreach, carer support, and services for populations that fall between departmental mandates. They are usually geographically concentrated, so their national coverage is thin even where their local quality is high, and they typically depend on a mixture of philanthropic funding, government grants and project money, which means their planning horizon is set by their funding cycle rather than by their patients' needs.

The partnership questions between such an organisation and a district service are concrete and are what a resident actually has to negotiate. Who accepts clinical responsibility for a person in

the organisation's care, and is that written down? What is the referral route in each direction, and what happens outside working hours? Who holds the clinical record, who prescribes, and who reviews? What happens to the people in the service if the grant ends? An arrangement with clear answers is a model; an arrangement without them is a friendship between two institutions, and it does not survive a change of personnel.

This section names no Indian organisation. That is a deliberate omission and worth stating plainly: naming a service is a factual claim about what that organisation does, and this chapter's sources do not carry a description of any individual Indian organisation. A resident preparing for an examination should know the sector's characteristic functions, which are set out above, and should read about specific organisations in their own published accounts rather than from a list learned second-hand.

25.3 The family as India's largest care system

In India the family is not an adjunct to care but the substrate on which everything else rests. Co-residence is common, supervision of medication is routine, transport to appointments is a family undertaking, and the family is usually the first to detect deterioration and the last line of support when nothing else is available. Reading that as a service model rather than as a cultural fact changes what a service is obliged to do about it.

If the family is a provider, then it has the same needs as any other provider. It needs to know what it is treating, which is information rather than instruction. It needs a route of access when things change, because a carer with no way to reach the service will present at an emergency department instead. It needs relief, since supervision without respite is not sustainable across the years these illnesses last. It needs financial support, because carer time is unpaid labour with an opportunity cost that falls disproportionately on women. And it needs somebody to notice its own health, since carer distress is both a consequence of the role and a determinant of whether the person being cared for stays at home.

Carer organisations and family self-help groups are the organised form of this. Their functions are advocacy, mutual support, and in some settings direct service provision, and they change the relationship between families and services from petitioning to negotiating. The techniques of family intervention itself, psychoeducation and the work on expressed emotion, belong to the notes cross-referred above and are not described here; what belongs here is the service obligation that follows from treating the family as part of the workforce.

TRAP **Describing family involvement as an Indian cultural advantage and stopping there.** It is an advantage, and it is also an unfunded transfer of long-term care to households, disproportionately to women, with no training, no relief and no income attached. Both halves belong in the answer. A candidate who writes only the first half has described a resource; a candidate who writes both has described a model with a cost.

25.4 Indigenous and traditional resources as a service model

Indigenous resources, the third limb, are what no other chapter treats as a service model, and they are difficult precisely because the phrase covers at least three different things. There are

traditional systems of medicine with their own texts, training and practitioners. There are religious and faith-based settings to which distressed people and their families go, sometimes in very large numbers. And there are local healers with no formal system at all, who are frequently the first point of contact and, in many districts, the most accessible one.

Treated as a service model, the honest ledger is short and should be stated without either romanticism or contempt. What these resources reliably supply is *access*, because they are local, affordable, culturally intelligible and available without a referral; *explanation*, in an idiom the family already uses; and *social sanction*, since consulting them carries none of the discredit a psychiatric attendance carries. What they do not supply is evidence of therapeutic effect for psychiatric disorder, and this chapter does not print one. The cultural and relational significance of indigenous healing, including the traditional teacher and disciple relationship read as a native model of the therapeutic bond, is taught in the Psychotherapies, clinical assessment, MSE and nosology notes, which places it at the cultural level and explicitly not as evidence-based treatment. This chapter prints no claim stronger than that.

The service consequence is a matter of position rather than of belief. These practitioners are already in the pathway, whether or not any service acknowledges them, so the operative question is not whether they should be, but what a formal service does about the fact. The design of that relationship, referral in both directions, what may safely be shared, what training is offered and what safeguards apply, is the subject of the traditional healers section of this chapter and is not settled here.

25.5 The one Indian community model with a trial behind it

Very few Indian community models have been evaluated with a controlled design, which is precisely why the ones that have been matter disproportionately. A community-champion model deserves attention on those grounds: in a cluster-randomised trial conducted in rural Gujarat, **the Atmiyata intervention had a significant effect on recovery from symptoms of depression and anxiety, with effects sustained at 8-month follow-up**.

What makes this instructive is the provider. The intervention is delivered by volunteer community champions, ordinary community members rather than health workers, which places it at the far end of the workforce spectrum from the specialist service. Three features of the design are worth extracting for any examination answer. The provider is drawn from the community being served, which addresses both access and acceptability at once. The design is cluster-randomised at village level, which is the right unit when the intervention acts on a community rather than on an individual. And the follow-up extends beyond the intervention period, which is what turns a demonstration of effect into a claim about durability.

IN INDIA

The decision this section supports is how to treat the non-governmental providers already working in your district. Treat them as a model rather than as goodwill: ask each one the four questions that make a model, who does what, on what money, with which referral route in each direction, and what they count. Where a volunteer or community-champion layer exists, the Indian evidence supports taking it seriously as a therapeutic component with a durable effect rather than as outreach, which means it deserves supervision, a defined scope and a route to escalate, exactly as a paid worker would.

25.6 What the Indian model literature is missing

The gap in this topic is not a gap in activity but a gap in evidence and in accounting, and it has a characteristic shape. There is no national inventory of voluntary-sector mental health provision, so nobody knows what exists or where. There is no measurement of what family care actually costs the households that supply it. Very few community models have been evaluated against any comparator, and those that have are concentrated in a handful of research settings, which raises the transportability question rather than settling it. And the most common form of publication is a description of a promising programme with no control condition and no follow-up, which is why so much of the field reads as encouraging and proves so little.

Because so much of the literature is descriptive, it is worth carrying a short appraisal routine for any Indian model-of-care paper, and it is the routine an examiner is testing when they ask what you think of a programme. Ask what the comparator was, and treat "before and after in the same people" as the weakest answer available, because a service that recruits people at their worst will improve without doing anything. Ask who was counted, and specifically whether the denominator is everybody offered the model or only those who stayed in it. Ask when the outcome was measured, since an assessment taken at the end of the intervention says nothing about durability. Ask who delivered it and whether that person exists outside the study, because a model delivered by a research team is not a model a district can run. And ask what it cost, since a model whose cost is unreported cannot be compared with anything the state is already funding.

Applying that routine to the Indian field produces a consistent verdict: a small number of models are properly evaluated, a large number are described, and the two groups are not the same models. That is not a criticism of the people running services, who are rarely funded to evaluate anything. It is an argument for building the evaluation question into the design of the next programme rather than attempting to reconstruct it afterwards.

None of that is a reason to dismiss the models. It is a reason to be precise about what is claimed for each, which is what the table below does.

MODEL	WHAT IT RELIABLY SUPPLIES	FUNDING AND GOVERNANCE	WHAT IS ESTABLISHED ABOUT EFFECT
Voluntary sector organisations	Long-term rehabilitation, day activity, vocational work, outreach to populations the public system misses	Philanthropy, grants and project money, with a planning horizon set by the funding cycle	Nothing at national level in this chapter's sources; no national inventory of provision exists
Family-based care	Supervision, accommodation, transport, early detection and the whole of long-term support	Unfunded, borne by households and disproportionately by women	Nothing quantified here; the cost to households is unmeasured
Community champions and volunteers	Local, acceptable, first-line psychosocial contact within the community itself	Programme or research funding, with volunteers rather than salaried staff	A cluster-randomised trial in rural Gujarat found improved recovery from depression and anxiety symptoms, sustained on follow-up
Indigenous and traditional resources	Access, explanation in the family's own idiom, and social sanction	Private, fee-for-service or offering-based, outside any health governance	No effect claim is printed here; the cultural reading is owned by another chapter and is explicitly not an evidence claim

PITFALL

Treating an unevaluated local programme as a model because it is admired. Enthusiasm, longevity and local reputation are not evidence, and a programme that has run for years without ever defining an output is not a model that anybody else can adopt. The kindest and most useful thing a visiting psychiatrist can do for such a programme is to help it define what it would count, because that is what converts it from a good thing into something a state could fund.

MNEMONIC

Provider, Purse, Pathway, Proof

The four questions that turn an arrangement into a model of care. Who does it, who pays for it, how does a person get in and out of it, and what would show that it worked. Ask them of a voluntary organisation, of a family and of a village volunteer alike, and the differences between the three limbs become clear immediately.

Volunteers
PROVIDER IN THE ONE INDIAN COMMUNITY MODEL WITH A CONTROLLED TRIAL HERE

Sustained
RECOVERY EFFECT AT FOLLOW-UP IN THAT TRIAL

Unfunded
BASIS ON WHICH FAMILY CARE IS SUPPLIED

No inventory
NATIONAL RECORD OF VOLUNTARY-SECTOR MENTAL HEALTH PROVISION

A model of care has a provider, a purse, a pathway and a proof; India's three great non-governmental models are the voluntary sector, the family and the indigenous practitioner, and of the three only a community-volunteer intervention carries a controlled Indian trial in this chapter's evidence.

Primary care, integration and service delivery models

26 · Integration of mental health into primary care and the role of medical officers

There will never be enough psychiatrists, and integration is the only answer anybody has that is not arithmetic nonsense. If mental disorder is common and specialists are few, then care has to be delivered by the part of the health system that is already everywhere and already sees everybody, which is primary care. Integration is the deliberate rebuilding of primary care so that it detects, treats, refers and follows up mental, neurological and substance-use conditions as ordinary business rather than as a specialist favour. This section is about how that is designed in India, who does what, and what is known about whether it is happening.

Two internal boundaries keep this section honest. The competencies, training and supervision that make non-specialist delivery possible, together with the decision-support instrument built for it, belong to the section of this chapter on task sharing. The escalation architecture of collaborative and stepped care, meaning who is added to a person's care and when, belongs to the section that owns those models. The national programme's own objectives and evolution belong to the section that owns the national programme, and are named here rather than restated.

26.1 What integration is, and three things it is not

Integration means that a mental health task is performed at the primary level by the staff who are there, using the same record, the same clinic session and the same supply chain as everything else they do. Being clear about what it is not is the faster route to the concept.

- **It is not a specialist clinic held in a primary-care building.** A visiting psychiatrist running a monthly session at a health centre is outreach. It is useful, it increases access, and it leaves the primary-care team's own practice exactly as it was, because nothing is done differently on the other twenty-something working days of the month.
- **It is not a screening programme.** Detection without a treatment response and a follow-up route generates identified and untreated people, which is worse than not looking, because it consumes trust as well as time.
- **It is not the transfer of all mental health care to primary care.** Integration implies a division of labour, and a division of labour implies that some conditions and some phases go upward. A model with no referral criteria has not integrated anything, it has abandoned people at the bottom of the system.

RULE **Integration is a two-way relationship or it is not integration.** Every design in this field has an upward limb, referral of what primary care cannot manage, and a downward limb, return of the person to primary care with a plan the primary-care team can execute. Systems build the upward limb because it is a request and neglect the downward limb because it is a gift. The result is a specialist service that fills with people who could be managed below it.

26.2 The medical officer's workflow, which is the heart of the item

Indian operational guidance is explicit about what the doctor at the primary level is expected to do. **At the primary and urban primary health centre level the medical officer's clinical workflow spans case detection, management, referral and follow-up of mental, neurological and substance-use disorders**. Those four verbs are the item in miniature, and each of them has a distinct requirement attached.

Detection requires that the presenting complaint be recognised for what it is, which in primary care is very often a bodily symptom, a sleep complaint, a request for a tonic, or a family member's description of behaviour rather than a patient's description of mood. It also requires time, which is the scarcest commodity in a busy general clinic.

Management requires a small formulary that is actually stocked, a small set of conditions the officer is expected to treat, and permission to treat them. Where any of the three is missing the workflow stops at detection.

Referral requires a named destination, a route and a reason to believe the person will be seen. A referral into a queue nobody can enter is a form of discharge.

Follow-up requires a record that persists between visits and somebody who notices non-attendance, which together are what the field calls **continuity of care**. This is the step most often absent, and its absence converts a chronic-disease service into a series of unrelated first visits.

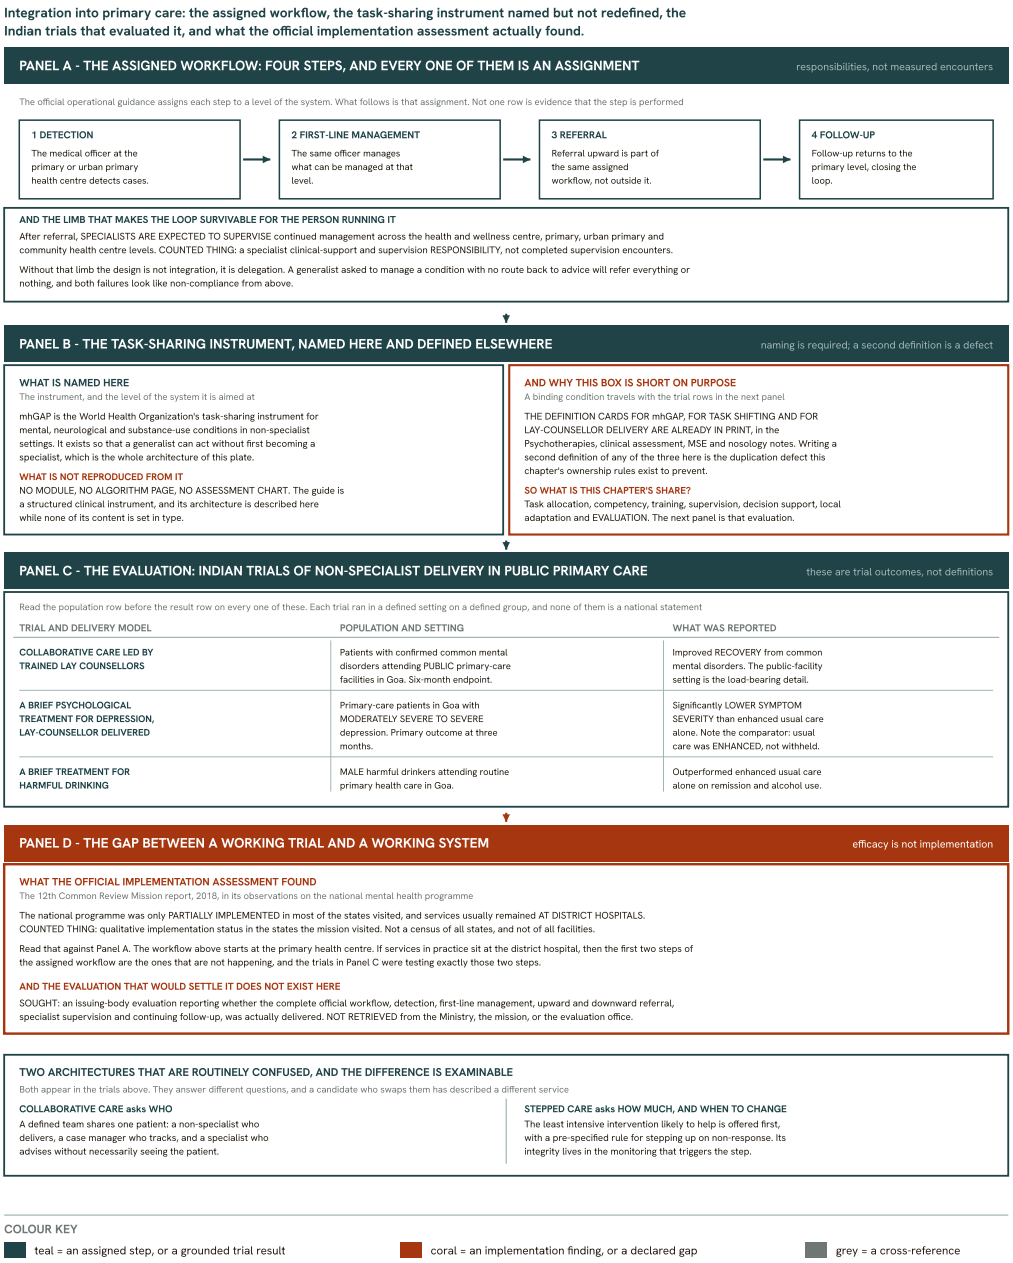


Fig 39.9 Primary-care integration: the assigned workflow, the instrument named, the trials that evaluated it, and the implementation finding that sits against all three. The medical officer at the primary or urban primary health centre is assigned case detection, management, referral and follow-up, and specialists are expected to supply continuing clinical support and supervision after referral, which is what separates integration from delegation. mhGAP is named here as the task-sharing instrument; its definition, and those of task shifting and lay-counsellor delivery, are already in print in the Psychotherapies, clinical assessment, MSE and nosology notes and are not rewritten. The evaluation is Indian and it is real: lay counsellor-led collaborative care improved recovery in public primary-care facilities, a brief lay-delivered treatment lowered depression symptom severity against enhanced usual care, and a counterpart for harmful drinking outperformed enhanced usual care alone. Against that, the official implementation assessment found the programme only partially implemented, with services usually remaining at district hospitals, and no issuing-body evaluation of the complete workflow was retrieved.

26.3 The upward limb, and the return leg that makes it a loop

The Indian programme design for severe illness is specific and worth learning as a sequence. **Within the national and district programme's service delivery, primary health centres refer people with severe mental disorders to the district hospital, and then follow up the treatment plan drawn by the psychiatrist there**. Read the second half of that sentence carefully, because it is the part that gets lost. The primary health centre is not simply a referring agency; it is the place where the specialist's plan is subsequently executed, month after month, which is where the great majority of the person's contact with the health system actually occurs.

Three practical requirements follow from the return leg and each is a common point of failure. The plan has to be written in a form the primary team can act on, which means named drug, dose, duration, what to watch for and when to send the person back. The medicine has to be available at the primary level, since a plan that requires a monthly journey to the district hospital for a supply has not returned the person anywhere. And the route back up has to be open without a fresh referral for a person who deteriorates, because a rigid one-way system punishes the primary team for accepting the patient in the first place.

TRAP **Describing integration as detection plus referral.** This is the commonest incomplete answer in the topic. Detection and referral are the first and third of four verbs in the official workflow, and a description that omits management at the primary level and follow-up after the specialist visit has described a triage service. The examinable content is the loop, not the arrow.

26.4 Supervision, which is what holds the whole design up

Supervision, meaning a continuing relationship in which a non-specialist can bring a real case to a specialist, is the component everything else rests on. A non-specialist asked to manage a condition they were not trained to manage will do it once, badly, and then stop, unless somebody is available to be asked. Indian guidance places that obligation explicitly: **after referral, specialists are expected to provide ongoing clinical support and supervision for the continued management of people with these conditions at the community health centre, primary health centre and health and wellness centre levels**.

Note what this is not. It is not a training event, which is a discrete input; it is a continuing relationship. Note also who it burdens: the same district psychiatrist who is already running the outpatient department is the person expected to supply it, which is why supervision is the component that quietly disappears when a post falls vacant. The design questions that decide whether it survives are mundane and answerable: what is the medium, is it scheduled or on demand, is it case-based or didactic, and is anybody's job description changed to include it. The competency framework and training architecture that sit behind supervision belong to the task-sharing section of this chapter.

IN INDIA

The decision this section supports is what to do with the referral letter you write as a district psychiatrist. Because the primary health centre is expected to follow the plan you draw, the letter is not a discharge summary but an operating instruction for a colleague who has no specialist training and no easy way to ask you a question. Write the drug, the dose, the intended duration, the two adverse effects that would matter, the specific circumstances that should bring the person back to you, and a way to reach you. A plan that assumes the primary team will interpret it is a plan that will be abandoned at the first uncertainty, and the person will simply stop attending.

26.5 What is actually happening, as opposed to what is designed

India has operational guidance, a defined workflow, a referral pathway and a supervision obligation. What it has less of is evidence that any of this is running. The most useful official observation available here is a review-mission finding: **the 2018 Common Review Mission found the national mental health programme implemented only partially by most of the states visited, with services usually available through district hospitals**.

That single sentence contains the central problem of this item. If services are usually available *through district hospitals*, then care has not descended to the primary level in most of the places that were looked at, which is the opposite of integration. Two cautions on scope: the finding describes states and facilities that a review mission visited rather than a census of the country, and it is qualitative, so it supports a direction and not a proportion.

The evaluation that would settle the question does not exist in this chapter's sources. No issuing-body evaluation reporting whether the complete official workflow, detection, first-line management by the medical officer, referral upward and downward, specialist supervision and continuing follow-up, has been implemented in Indian primary-care facilities was retrieved. It was sought from the health ministry, the national health mission, the national health systems resource centre and the government's evaluation office, and what was found was either an earlier review, a document predating the current operational guidance, or a procurement notice for an evaluation rather than its findings.

Four	District hospital	Partial	Not retrieved
VERBS IN THE OFFICIAL PRIMARY-CARE WORKFLOW: DETECT, MANAGE, REFER, FOLLOW UP	WHERE SERVICES WERE USUALLY AVAILABLE IN THE STATES A REVIEW MISSION VISITED	IMPLEMENTATION FOUND IN MOST STATES VISITED, QUALITATIVE AND NOT A CENSUS	ANY EVALUATION OF THE COMPLETE WORKFLOW IN INDIAN PRIMARY CARE

26.6 Where integration breaks, in the order it breaks

The failure points are predictable, and a resident who can name them in sequence can both diagnose a local service and answer the question.

STEP	WHAT HAS TO BE TRUE	HOW IT FAILS	WHAT WOULD SHOW IT
Detection	Time in the consultation, and recognition of presentations that are not complaints of mood	The presentation is treated as the physical complaint it arrived as, and nothing further happens	Recorded diagnoses at the primary level, against expected frequency
Management	A stocked formulary, a defined scope of practice and permission to prescribe within it	Detection without treatment, so the officer stops detecting	Drug availability at the primary level, and prescriptions actually issued there
Referral upward	A named destination that will see the person	Referral into an inaccessible queue, which the family experiences as refusal	Proportion of referrals that result in a specialist contact
Return downward	An executable plan and the medicines to execute it	The person keeps returning to the district hospital for routine review, which silts up the specialist clinic	Proportion of stable people reviewed at the primary level rather than at the hospital
Supervision	A scheduled, case-based relationship with somebody's job description changed to include it	It becomes voluntary, then occasional, then nothing, especially when a post is vacant	Scheduled supervision contacts held, not training events delivered
Follow-up	A persisting record and somebody who notices non-attendance	Every visit becomes a first visit and continuity is lost	Retention in care over a defined period, with the denominator stated

PITFALL

Judging a primary-care mental health service by how many people it detected. Detection is the cheapest number to produce and the easiest to increase, and on its own it says nothing about whether anybody was treated or stayed in treatment. The number that reflects an integrated service is retention: how many of the people started on treatment at the primary level are still in care after a defined period, with the denominator stated. If a service cannot produce that, it does not yet know whether it works.

Reading the table upward rather than downward gives the other examinable lesson. A service that is failing at follow-up cannot be fixed by more training in detection, and a service failing at management because nothing is stocked cannot be fixed by supervision. Each step depends on the ones before it, so the correct first question about any underperforming integrated service is which step is the earliest one that is broken. Programmes that repeat detection training every year while the formulary remains empty are extremely common, and they are not failing through incompetence; they are fixing the step that is easiest to fix.

26.7 The workforce question, and why this section prints no benchmark

The argument for integration is usually opened with a workforce statistic: so many psychiatrists for so many people, against a recommended minimum. This section prints no such benchmark, and the reason is worth teaching because it is a lesson about evidence rather than about workforce.

Several psychiatrist-per-population figures circulate in Indian teaching, they differ from one another, and the recommended minimum they are usually measured against cannot be traced to an issuing body in this chapter's sources. A ratio quoted without a source, a reference year and a stated definition of who counts as a psychiatrist is not a fact; and a minimum attributed to an international body that did not issue it is a false attribution, which is a more serious error than a wrong number because it borrows authority as well as data. The defensible position is to make the argument structurally, which loses nothing: specialists are few and unevenly distributed, mental disorder is common and present in every district, and no plausible expansion of specialist training changes that arithmetic within a working lifetime. That argument does not need a ratio, and it cannot be wrong.

GOOD TO REMEMBER

When a primary-care colleague telephones about a patient, the useful reply is not a diagnosis but a decision they can act on today: what to start or change, what to watch for, when to review, and what should bring the person back to you. Answering at that level takes less of your time than a referral does, keeps the person near home, and is the single most effective piece of integration available to an individual psychiatrist regardless of what the programme is doing.

Integration is a loop, not an arrow: detect, manage, refer upward and take the person back with a plan the primary team can execute, held together by continuing specialist supervision, and Indian review findings indicate that services in most states visited still sat at the district hospital rather than below it.

27 · WHO mhGAP intervention guide and task-shifting

If the specialist workforce cannot grow fast enough, the only remaining variable is who is allowed to do what. Task sharing is the deliberate redistribution of defined clinical tasks to workers with less training, supported by structured decision aids, supervision and a clear escalation route. The World Health Organization's mhGAP Intervention Guide is the instrument built to make that redistribution safe for mental, neurological and substance-use conditions. This section treats both as design problems: how tasks are allocated, what competency means, how training and supervision are constructed, what a decision aid actually does, how much a country may change, and what the evidence shows.

The definitions themselves are owned elsewhere and are named rather than rebuilt. Task-shifting and task sharing as constructs, the mhGAP Intervention Guide as a document, and lay-counsellor-delivered treatment as a category are all defined in the Psychotherapies, clinical

assessment, MSE and nosology notes, which carries them with their evidence base. The primary-care workflow into which any of this is delivered belongs to the section of this chapter on integration, and the escalation architecture of stepped care to the section that owns it. No module, algorithm or assessment chart from the guide is reproduced anywhere in this chapter.

27.1 Task allocation is a design decision, and it is made task by task

The question is never whether a cadre is "capable". It is whether a specific task, defined narrowly enough to be trained and checked, can be performed by a specified worker under specified conditions with an acceptable failure rate. Four properties make a task a good candidate for redistribution, and the same four explain why some tasks resist it.

- **The task is bounded.** It can be described as a procedure with a start, a finish and a small number of decision points, rather than as clinical judgment applied to an undefined situation.
- **The decision rule is explicit.** Somebody can write down what to do in each case that arises, which is precisely what a decision aid is for.
- **The failure is visible.** If the task is done badly, somebody notices, either because the person deteriorates in a recognisable way or because a check exists.
- **An escalation route exists.** The worker can hand the situation upward quickly when it exceeds the boundary, and doing so is treated as correct practice rather than as failure.

Applying those four in order settles most disputes. Delivering a manualised brief psychological treatment for a defined condition scores well on all four. Adjusting an antipsychotic in a person with a complex physical comorbidity scores badly on the first two. Recognising that somebody is unwell and needs to be seen scores well on all four and is therefore the task most widely and most safely redistributed anywhere in the world.

RULE **Task sharing redistributes a task with its supports, or it is simply understaffing.** The package is training plus a structured decision aid plus supervision plus an escalation route. Remove any one of the four and the model becomes a less qualified person doing a harder job alone, which is not what any of the evidence tested.

27.2 Competency rather than cadre

Competency, meaning what a worker can be observed to do rather than what they hold a certificate in, is the organising idea of this whole section. The professional reflex is to ask what qualification somebody holds. The task-sharing question is what they can demonstrably do, which is a different and more useful thing to know. The practical consequence is that programmes specify competencies, then select whoever can acquire them, which is how community health workers, nurses, counsellors without a psychology degree and general medical officers all end up delivering parts of the same programme in different settings.

Selection matters more than it looks. Programmes that recruit for attributes the training cannot supply, literacy sufficient for the materials, standing in the community, the ability to hold a conversation about distress without flinching, and a willingness to be supervised, achieve competency faster than programmes that recruit on formal qualification alone. And competency

has to be assessed rather than assumed: attendance at a training course is a record of presence, and observed performance on the task is a record of competence.

27.3 Training, and why most of it decays

Training in this field fails in a characteristic way. A cascade model, in which a small group is trained centrally and each trains the next tier, loses content at every step; a single workshop produces a measurable rise in knowledge that decays within months; and knowledge gained in a classroom transfers poorly to a clinic without practice under observation.

What survives has a recognisable shape: it is short, repeated rather than single, built around cases rather than around lectures, includes observed practice, and is followed by continuing contact with the trainer. That last element is the bridge from training into supervision, and it is why the two should be commissioned as one thing rather than as a course and a hope.

27.4 Supervision, the component that is always cut first

Supervision in a task-sharing programme is not the periodic review of a subordinate. It is the mechanism by which a bounded task stays bounded: the point at which cases that have drifted outside the worker's remit are identified and moved, at which errors are corrected before they become habits, and at which the worker is retained, because isolated non-specialists doing emotionally demanding work leave.

Its design parameters are few and are worth being able to state: frequency, who supervises, whether it is individual or group, whether it is case-based or general, and whether attendance is protected time or personal generosity. Group case-based supervision at a defined interval, delivered by somebody one step up rather than by the most senior person available, is the arrangement that most often survives contact with a real service, because it is affordable and because it uses the tier immediately above rather than the scarcest tier.

TRAP **Presenting task sharing as a solution to a shortage of specialists.** It redistributes specialist time rather than removing the need for it. Every task-sharing model consumes specialist time in training, in supervision and in receiving escalated cases, and a district with no specialist at all cannot run one properly. The correct framing is that task sharing multiplies the reach of the specialist who exists, not that it substitutes for one who does not.

27.5 What a decision-support instrument does

The **mhGAP** Intervention Guide is the instrument the field is built around, and its function is worth describing precisely because the description is what an examiner wants. Architecturally it is organised by presentation rather than by diagnosis, it converts assessment and management into stepwise sequences a non-specialist can follow, it specifies what must be excluded before proceeding, and it states explicitly at which points the person should be referred or discussed with a specialist. In other words it externalises the parts of clinical reasoning that ordinarily live in a trained clinician's head, so that the worker's judgment is required in fewer places and those places are marked.

Its limits follow from the same design. A presentation-based sequence handles the common and the typical and is not built for the atypical, the comorbid or the treatment-resistant, which is exactly why the referral points are part of the instrument rather than an afterthought. And a guide is not a programme: possessing it changes nothing without training, supervision and a service to escalate into. This chapter names the guide, describes what it does, and reproduces no part of its content.

27.6 Local adaptation, and what may not be changed

Every country adapts, and adaptation is expected rather than tolerated. The distinction to hold is between adapting the *surface* and adapting the *structure*. **Surface adaptation** covers language, the idiom in which symptoms are described, the examples used, the cadre selected, the setting, and alignment with the national formulary and referral map, and all of it improves fidelity to intent. **Structural adaptation**, meaning removing a referral trigger, widening the task boundary, dropping a step because it is inconvenient, or omitting supervision because there is nobody to supply it, changes the intervention into something the evidence does not cover.

The examinable formulation is that adaptation should preserve the safety architecture and change everything else. A programme that has translated the materials, chosen a locally appropriate cadre and aligned the drugs with what the district actually stocks has adapted well. A programme that has kept the materials in their original form and quietly abandoned supervision has adapted badly while appearing faithful.

27.7 The Indian evidence, which is unusually strong

India contributed two of the most-cited randomised evaluations in this field, both testing brief psychological treatments delivered by lay counsellors in routine primary care against enhanced usual care. For depression, **participants receiving the Healthy Activity Program in addition to enhanced usual care had significantly lower symptom severity than those receiving enhanced usual care alone**, in primary-care patients with moderately severe to severe depression. For harmful drinking, **Counselling for Alcohol Problems delivered by lay counsellors in addition to enhanced usual care was better than enhanced usual care alone for harmful drinkers in routine primary health-care settings**.

Four features make this pair worth learning rather than merely citing. The comparator in both is enhanced usual care rather than nothing, so the finding is not that attention beats neglect. The setting is routine primary care rather than a research clinic, which is the setting the result has to generalise to. The provider is a lay counsellor, which is the whole point of the design. And the two conditions are different, one a mood disorder and one a substance-use problem, which suggests the approach is a method rather than a treatment for one condition.

The proper use of these trials in an answer is as evidence about *delivery*. They establish that non-specialists, trained and supervised, can deliver a manualised psychological treatment in an Indian primary-care setting to a standard that outperforms the enhanced usual care they were compared against. They do not establish that any untrained worker can deliver anything, and they say

nothing about what happens to such a programme once the trial's training and supervision structure is withdrawn.

IN INDIA

Because the delivery question has been answered here rather than borrowed, the decision facing an Indian service is not whether non-specialist delivery can work but whether the district can supply the supervision that made it work. That converts the planning question into an arithmetic one a resident can actually do: how many workers, at what supervision frequency, in groups of what size, adds up to how many hours of somebody's week, and whose week is it. A programme that cannot answer that has not been designed, and it will reproduce the trials' structure for as long as its initial enthusiasm lasts.

One further caution belongs with the Indian evidence and is the kind of point that separates a good answer from a fluent one. Trials of this sort recruit, train and supervise their workers to a standard that a research team can guarantee and a district usually cannot, so the trial estimates what the model can do under favourable conditions rather than what it will do under ordinary ones. That is not a criticism of the trials, which were conducted in routine settings precisely to narrow that gap. It is a reminder that the implementation question, whether the supervision structure survives once the researchers leave, is a separate empirical question from the efficacy question, and that this chapter's sources answer the first and not the second.

27.8 The design decisions, assembled

DECISION	THE QUESTION TO ANSWER	THE COMMON FAILURE	WHAT IT COSTS IF YOU GET IT RIGHT
Which task	Is it bounded, rule-governed, checkable, and escalatable	Redistributing judgment rather than a procedure	Time spent writing the boundary down before anybody is trained
Which worker	What competencies are required, and who can acquire them	Selecting on qualification rather than on demonstrable competency	A selection process, and the willingness to reject candidates
What training	Short, repeated, case-based, with observed practice	A single workshop, cascaded through tiers, measured by attendance	Repeated contact rather than one event
What supervision	Frequency, supervisor tier, group or individual, protected or voluntary	Dropped when a post falls vacant, and never restored	A standing claim on the time of the tier above the worker
What decision support	Presentation-based sequences with explicit exclusions and referral points	Distributing the document and calling it implementation	Adaptation work, and alignment with the local formulary
What adaptation	Change the surface, preserve the safety architecture	Preserving the materials and abandoning the supervision	Explicitly recording what was changed and why
What evaluation	Outcome in the person, not activity in the worker	Reporting numbers trained instead of people improved	An outcome measure collected routinely from the start

PITFALL

Counting workers trained as the programme's result. Trained workers are an input, and they decay: people leave, get redeployed, or stop doing the task when nobody asks about it. A programme reporting cumulative numbers trained across several years is reporting a quantity that has no relationship to how many people are delivering the task today. The number that matters is how many trained workers are currently active and supervised, and almost nobody publishes it.

MNEMONIC

Task, Trainee, Training, Tools, Talk-to-someone

The five things a task-sharing programme has to specify: the bounded task, the worker who will do it, the training that produces the competency, the decision-support tools that hold the reasoning, and the supervision and escalation route. The last is the one that disappears, and it is the one all the evidence depended on.

Bounded	Enhanced usual care	Routine primary care	Supervision
THE PROPERTY THAT MAKES A TASK REDISTRIBUTABLE	THE COMPARATOR IN BOTH INDIAN LAY-COUNSELLOR TRIALS	THE SETTING THOSE TRIALS RAN IN, NOT A RESEARCH CLINIC	THE COMPONENT CUT FIRST AND THE ONE THE EVIDENCE RESTED ON

Task sharing redistributes a bounded task together with its training, its decision support, its supervision and its escalation route, and the two Indian lay-counsellor trials that outperformed enhanced usual care in routine primary care are evidence about that whole package rather than about the counsellor alone.

28 · Telepsychiatry: models, telepsychiatry guidelines and Indian regulatory framework

Telepsychiatry earns its marks as a service, not as a gadget. The examinable question is never whether a psychiatrist can be seen on a screen; it is what a health system gains, and what it must build, when the specialist and the patient stop having to be in the same building. That reframing decides everything downstream: who is stationed where, who is permitted to start the encounter, what is counted, and what the counting is allowed to prove. Candidates lose the marks in three predictable places. They describe the modalities of remote consultation, which is a technology answer to a service question. They recite a national helpline total as though it were a number of patients. And they attribute a governance position to an instrument whose current status they have not checked. The professional-conduct frame around remote practice, the consent and confidentiality obligations at a distance, and the identity of the operative regulatory code all belong to the Forensic ethics, consent and legal procedure notes, and are not restated here. What follows is the delivery architecture, the national programme, the monitoring discipline and the equity argument.

28.1 The service problem, stated before any solution

A specialist mental-health workforce is small, unevenly distributed and slow to grow. Every scheme for extending it is a variation on the same manoeuvre: move the scarce expertise, or move the decision that needs it, so that the specialist's time is spent only on what nobody else can do. Referral moves the patient. Outreach moves the specialist. **Task-sharing** moves the task. **Remote consultation** moves the encounter itself, and it is the only one of these that does not consume travel time at either end.

That is the whole economic case, and it explains why the design questions are the ones they are. If the constraint is specialist time, a remote service that simply converts face-to-face appointments into video appointments has bought convenience and no coverage: the same clinician sees the same number of people, and the queue behind them is unchanged. Coverage only rises when the remote link lets someone who is *not* a psychiatrist hold the patient contact while the psychiatrist supplies the decision. A telepsychiatry programme that has not answered who holds the contact has not yet been designed.

RULE **A telepsychiatry programme is judged by the population it can now reach and by whose time it saves, never by the technology it deploys.** Bandwidth, platform and picture quality are preconditions. They are not the intervention, and an appraisal that stops at them has appraised nothing that could change a coverage figure.

28.2 Hub and spoke, and what the architecture is actually doing

The dominant service design for remote specialist care is the **hub-and-spoke model**. A small number of **hubs** concentrate the scarce resource, in this case psychiatrists and the supervisory and record-keeping apparatus around them. A much larger number of **spokes** sit where the population is, staffed by whoever the health system already has at that level. The link between them carries clinical decisions in one direction and clinical information in the other.

Read structurally, the model does three separable things, and a programme can succeed at one while failing at the others. It *concentrates* expertise, so that a psychiatrist is not the rate-limiting step in every district simultaneously. It *distributes* access, so that the distance a patient travels falls to the nearest spoke rather than the nearest specialist. And it *standardises*, because a hub that supervises many spokes tends to impose one assessment format, one prescribing habit and one escalation route across all of them. The third effect is the one most often overlooked and is frequently the largest clinical gain, because it raises the floor of care at the periphery rather than the ceiling at the centre.

The corresponding failure modes are equally separable. A hub with too many spokes becomes a new bottleneck with a longer waiting list. A spoke with no trained staff becomes a video booth, which returns the model to moving the patient. And a link with no agreed **escalation route**, meaning a pre-arranged answer to who attends an acute presentation at the spoke and where the person goes, leaves the periphery holding risk it was never resourced to hold. That is where remote services generate their most serious incidents.

This chapter cannot state the official Indian architecture in the terms its own programme documents use. What was sought was a directly retrieved official passage defining the hub-and-spoke or tiered service structure for Indian telepsychiatry, including the roles of peripheral cells and specialist hubs; the programme and institutional sources approached either failed to open, exceeded retrieval limits or returned server errors, and no such passage was obtained. The reasoning above is service-design reasoning and is offered as such. An answer that describes the generic architecture and then says plainly that the official Indian tiering is not established from a source it has read is a stronger answer than one that supplies a confident tier list.

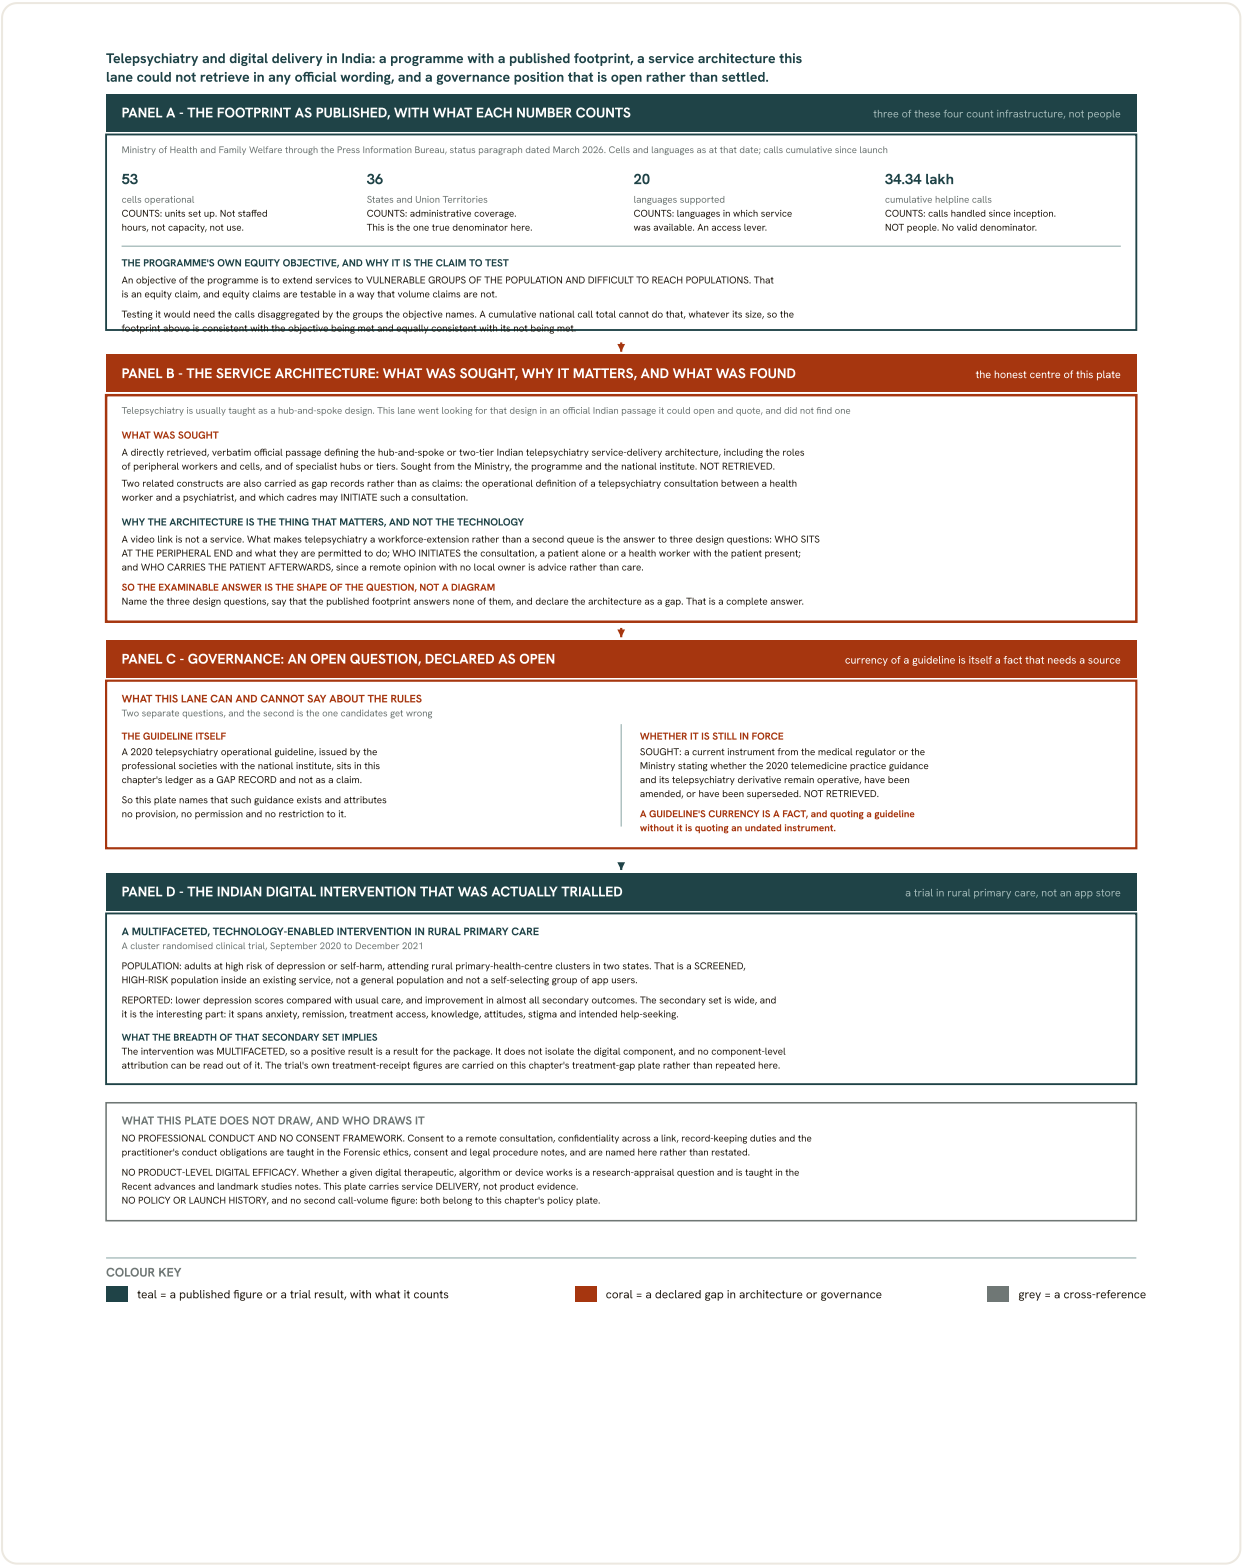


Fig 39.10 Telepsychiatry and digital delivery: a published footprint over an architecture that is a declared gap. The programme reports 53 cells across 36 States and Union Territories with support in 20 languages, and more than 34.34 lakh cumulative helpline calls; three of those four count infrastructure and the fourth counts calls rather than people. Its stated objective is to extend services to vulnerable and difficult-to-reach populations, which the footprint can neither confirm nor refute because it is not disaggregated. The service architecture itself is the honest centre of the plate: a verbatim official passage defining the hub-and-spoke or two-tier design, and the roles of peripheral workers and specialist hubs, was sought and not retrieved, and no current instrument confirming whether the 2020 telemedicine guidance remains operative was retrieved either. The one Indian trial drawn here reported lower depression scores and improvement in almost all secondary outcomes in a high-risk rural primary-care population, for a multifaceted package that does not isolate its digital component.

MNEMONIC

Place, Person, Permission, Proof

The four axes on which any remote mental-health programme is specified, and the order in which they fail. Where are the spokes and where is the hub. Who staffs each end. Who is permitted to initiate the encounter. What is measured well enough to show whether the reach was real.

28.3 The national instance, and what it was set up to do

India's national tele-mental-health service is delivered through **Tele-MANAS**, and its own statement of purpose is the right place to begin, because a programme's declared objectives are what its monitoring should be reporting against. Among those objectives is the explicit commitment to extend services to vulnerable groups of the population and to populations that are difficult to reach. That sentence is worth reading as a design instruction rather than as a preamble: it makes equity of reach a stated aim of the programme, which means underserved reach is a legitimate thing to hold the programme to, not an external criticism of it.

The programme's published national status figures describe its footprint. At its most recent published national status update, reported by the Ministry of Health and Family Welfare, there were **53** operational cells across **36** States and Union Territories, with service available in **20** languages, and **more than 34.34 lakh** calls received and addressed through the helpline since the service began.

53	36	20	more than 34.34 lakh
OPERATIONAL CELLS	STATES AND UNION TERRITORIES COVERED	LANGUAGES OF SERVICE	CALLS ADDRESSED SINCE INCEPTION, NOT PEOPLE

Each of those quantities counts a different thing, and the differences are the examinable content. A cell is a staffed unit, so the cell count is a supply figure. Coverage of States and Union Territories is an administrative reach figure and says nothing about depth within any one of them. The language count is an access figure and is arguably the most under-discussed of the four, because a helpline in a language the caller does not think in is a helpline the caller does not use. The call total is a **cumulative activity figure** and is the one that is misread most often, because activity is not the same object as the **unique users** a coverage claim would need.

28.4 Reading a programme's own numbers without laundering them

A national figure is only a fact when four things travel with it. Who issued it, what period it refers to, which population it describes, and what it is a proportion of. Strip any one and the number becomes a rumour with a decimal point in it. Applied to a helpline total, the fourth question is fatal: a cumulative count of calls has no valid denominator at all. It is not divided by the population, because it is not a count of people; it is not divided by callers, because one caller may call many times; and it cannot be compared with a prevalence, because the two count incommensurable things.

WHAT THE FIGURE IS	WHAT IT CAN SUPPORT	WHAT IT CANNOT SUPPORT
Number of operational cells	A statement about staffed service capacity	Any claim about how much of that capacity was used
States and Union Territories covered	Administrative reach across the country	Depth of reach inside any one of them
Languages of service	A concrete access provision for callers	That every caller reached service in their own language
Cumulative calls addressed	Total activity handled since the service opened	Number of individuals helped, or any rate per population

TRAP **Reading a cumulative helpline call total as a count of people served, and then dividing it by a population to produce a coverage rate.** Candidates do this because the number is large, official and easy to remember, and because every other figure in a public-health answer is a rate. The correct sentence names the count as calls received and addressed since the service began and stops there. If an examiner wants a coverage claim, the honest answer is that the programme publishes activity, not unique users, so a coverage rate cannot be computed from it.

IN INDIA

Because reaching vulnerable and difficult-to-reach populations is a stated objective of the national tele-mental-health programme, a state cell planning its rota has a specific and answerable design question rather than a general aspiration: which of the populations named in that objective is it currently failing to receive, and what in the service is producing the failure. Language of service, hours of operation and the referral destination offered at the end of the call are the three levers that can be moved locally, and the objective makes each of them a monitoring question rather than a courtesy.

28.5 The regulatory frame, and the currency question inside it

Remote practice in India operates inside a national telemedicine framework and a professional-conduct code, and this chapter names them as the governing frame and goes no further. The provisions themselves, the regulation numbers, the sequence by which the newer conduct regulations were notified and their subsequent status, and the identity of the code that is operative today are taught in the Forensic ethics, consent and legal procedure notes, together with consent to the modality, identity verification at both ends, confidentiality at a distance and the situations in which a remote encounter is not enough. Repeating any of that here would produce a second and possibly divergent account of a statutory position, which is the single most damaging thing a service chapter can do to a legal one.

Two regulatory facts a service chapter does need, and cannot supply, are declared here rather than filled in from recollection. The first is the identity, issuing body and year of the dedicated Indian telepsychiatry operational guidance that practitioners commonly refer to; what was sought was that document itself, and it was not established from a source this chapter opened. The second is more consequential and is a genuine live uncertainty: no current national instrument was retrieved that states explicitly whether the standing telemedicine practice

guidance and the telepsychiatry guidance derived from it remain operative, have been amended, or have been superseded. Multiple official routes were attempted and each failed to open or returned unusable content.

The teaching point survives the gap, and is arguably sharpened by it. Guidance for remote practice has changed repeatedly and quickly in every jurisdiction that has issued any, so the currency of the instrument is part of the clinical question and not an administrative footnote. A candidate who names the framework, states that its current standing must be checked against the issuing body rather than assumed, and points to the chapter that owns the conduct code has answered correctly. A candidate who recites a version and a year from memory has asserted the one thing most likely to have moved.

28.6 Who may open the encounter, and why that answer sets the ceiling on coverage

Return to the architecture. The hub extends the workforce only if a non-specialist may hold the patient and initiate the contact with the specialist. Everything about the size of a programme therefore depends on a permission: whether a nurse, an auxiliary nurse midwife, an allied health professional or another cadre of health worker may start a telepsychiatry consultation, and under what supervision. If only a registered medical practitioner may initiate, the spokes must be staffed by doctors and the model inherits the shortage it was built to relieve. If a trained health worker may initiate, the spoke can sit at the level where the population actually presents.

The specific Indian operational position on this is a declared gap in this chapter. Two related propositions were sought and not established from an opened source: an operational definition of telepsychiatry service delivery conducted between a health worker and a psychiatrist, and an operational statement permitting nurses, allied health professionals, auxiliary nurse midwives and other health workers to initiate such a consultation. Both are exactly the kind of proposition a candidate is likely to half-remember from a slide, and both change the shape of a service completely depending on which way they read.

What can be taught safely, and what an examiner is really testing, is the consequence structure. The permission determines the staffing of the spoke. The staffing determines the training burden. The training burden determines how fast the programme can scale and how uniform its quality will be. And the same permission determines the medico-legal question of who is responsible for the patient during and after the link, which is where the conduct code becomes unavoidable and where the reader is sent to the chapter that owns it.

28.7 Governance, monitoring and the evaluation a remote service owes

Remote delivery makes some things easier to measure and some things easier to lose. It is easier to count contacts, because they pass through a system that logs them. It is harder to know what happened afterwards, because the person who would have noticed is at the other end of a link that has closed. Monitoring has to be designed against that asymmetry rather than around the data that happen to accumulate.

A remote mental-health service should be able to state the following about itself, and a service that cannot is being appraised on its activity log alone.

- Who was reached, described by the populations the programme says it exists to reach, and not only by total volume.
- What proportion of contacts ended with a named destination, meaning a specific onward service, review or clinician, rather than advice alone.
- Whether the destination was reached, which requires a follow-up mechanism the link itself does not provide.
- How acute risk identified during a remote contact was escalated, to whom, and with what agreed local response.
- What the spokes needed and did not get, which is the only signal that reliably distinguishes an under-trained spoke from an under-used one.

The Indian reach argument for remote specialist care, including the contribution of the non-governmental sector to extending it, is developed in the Schizophrenia: management, antipsychotics and adverse effects notes, and is not re-derived here.

PITFALL

Designing the escalation route last. A remote service is commissioned around the consultation, and the acute pathway is added afterwards if at all, so the spoke discovers during its first emergency that nobody has agreed who attends, where the person is taken, or which clinician holds responsibility until they arrive. The escalation route is not an appendix to a telepsychiatry protocol; it is the part that determines whether the service can safely accept unfiltered presentations at all.

28.8 Equity, and the way a remote service quietly reproduces the gap it was built to close

Remote delivery is often presented as an equity intervention, and it can be one, but only against an explicit account of who is excluded. Every remote service imposes a set of entry requirements that face-to-face care does not: a device, a connection, sufficient privacy at the patient's end to speak freely, enough literacy or confidence to navigate a call, and a language the service offers. Each of those requirements is unevenly distributed in the same direction as the underlying disadvantage, so an unexamined remote service can expand access fastest among those who already had the most.

This is why the programme's own equity objective matters operationally. If reaching vulnerable and difficult-to-reach populations is a declared aim, then the entry requirements above become the programme's business, and the counter-measures are concrete: service in more languages, a staffed spoke where the person has no private device, hours that fit shift and agricultural work, and a call destination that exists in the caller's own district. None of that follows automatically from putting a service online, and the difference between a remote service that narrows a gap and one that widens it lies almost entirely in whether those counter-measures were built.

Telepsychiatry extends a workforce only when someone other than the specialist may hold the patient; everything else about the model, including its numbers, follows from that permission.

29 · Collaborative and stepped-care models of mental health delivery

Most people with a common mental disorder are already in contact with a health service; almost none of them are in contact with a psychiatrist. Collaborative and stepped care are the design answers to that arithmetic. They are not treatments, and they are not referral pathways with better manners. They are ways of organising a primary-care system so that the decisions requiring specialist judgement reach a specialist while the contact, the follow-up and most of the delivery stay where the patient already is. The marks here sit on structure rather than on content: what the roles are, what triggers a change of treatment, where the joins are, and what the system must measure to know whether any of it worked. Case management and assertive community treatment as service models, with their caseload logic and target populations, belong to the Schizophrenia: management, antipsychotics and adverse effects notes; the general-hospital liaison models and the classic economic argument for funding a liaison service belong to the Consultation-liaison, geriatric and transcultural psychiatry notes, which also records that the Indian cost-effectiveness evidence for such services is what is missing. Neither is restated here.

29.1 Stepped care as a rationing principle, stated honestly

Stepped care organises an intervention set by intensity and offers the least intensive intervention that is likely to work, with an explicit rule for moving up when it does not. Two commitments make it work and both are frequently dropped in practice. The first is that the lowest step must be a real intervention with a real chance of benefit, not a holding position. The second is that the rule for stepping up must be written down before the patient arrives, applied on a schedule, and triggered by an observation rather than by the patient's persistence or the clinician's patience.

It is worth being candid that stepped care is a rationing principle. A system adopts it because it cannot give everybody the most intensive option, and it distributes the scarcest resource by escalating only those who have already failed something cheaper. That is a defensible allocation, and it is defensible precisely to the extent that the escalation actually happens. Where the stepping rule is absent, stepped care degenerates into a queue with a therapeutic vocabulary, and the people who do worst are the ones least able to keep coming back to ask.

A distinction worth fixing early is between the lowest step and no step at all. Watchful waiting, in which a clinician defers treatment and reviews later, is a legitimate clinical decision but it is not the bottom rung of a stepped-care system unless the review is scheduled and somebody owns it. The bottom rung of a working system is an active low-intensity intervention with a named deliverer, a defined content and a defined end point. The difference is not semantic: in the first arrangement the patient returns if they deteriorate, and in the second the service comes back to look.

There is also a stratification logic that stepped care is sometimes confused with. Stratified allocation assigns the patient at the outset to the intensity their presentation appears to warrant, and does not require failure at a lower level first. Stepped care allocates by observed response rather than by initial appearance. Most real programmes are hybrids, and it is worth saying so rather than pretending purity: severity, risk and previous treatment history are used to enter people at a higher step, while response governs movement thereafter. What must not happen is that the hybrid becomes an excuse to hold everybody at the bottom.

RULE **A stepped-care system is defined by its stepping rule, not by its steps.** Any service can list interventions in order of intensity. What makes the arrangement stepped care is a pre-specified observation, taken at a pre-specified interval, that obliges the system to change what it is doing. Without that, the arrangement is a menu.

29.2 Collaborative care, and the three roles that constitute it

Collaborative care is a specific organisational form and not a synonym for good communication. Its defining feature is that a named person, distinct from both the treating primary-care clinician and the specialist, takes responsibility for keeping a defined group of patients under active review. That role is usually called the **care manager**, and the model stands or falls on whether it exists as a funded post rather than as an aspiration attached to somebody's existing job.

Three roles interlock. The primary-care clinician retains clinical responsibility, prescribes and remains the patient's medical home. The care manager holds the register, makes the scheduled contacts, delivers or supports the low-intensity intervention, records the agreed measure and raises the cases that are not improving. The consulting specialist reviews those cases, usually without seeing most of the patients, and advises on the ones that are stuck. The specialist's time is therefore spent on decisions rather than on encounters, which is the whole point: one psychiatrist supervising a caseload through a care manager can influence a population many times larger than the one they could see.

Two structural features are easy to miss and are frequently examined. The first is the **population register**: collaborative care is delivered to an enumerated list of patients, not to whoever returns. A patient who stops attending is a case the register flags, whereas in ordinary care they are simply absent. The second is that the specialist review is **caseload-based** rather than referral-based. The specialist looks systematically at those who are not improving, which inverts the usual selection, where the cases that reach a specialist are the ones somebody remembered to refer.

MNEMONIC

Register, Role, Review, Route

The spine of any collaborative-care arrangement. An enumerated population register. A named care-manager role that is somebody's actual job. A scheduled review carrying an agreed measure. And a route that obliges the system to do something different when the measure does not move.

29.3 Measurement-based escalation, and what a measure is for here

The measurement in measurement-based care is not for diagnosis and not for research. It exists to make deterioration and non-response visible to a system that would otherwise notice neither. That purpose sets its properties: it must be short enough to be repeated at every contact, stable enough that a change means something, and administered by whoever is actually in the room, which in this model is the care manager and not the specialist.

The clinically important discipline is that the measure must be attached in advance to a consequence. A rating repeated at intervals with no pre-agreed action attached is documentation, and documentation absorbs the same clinical time as a decision while producing none. The service therefore has to specify, before the programme opens, what change counts as adequate progress, at which review the question is asked, and what happens on each answer. Naming and describing an instrument is legitimate teaching; this chapter typesets no instrument's items, response options, scoring rule or thresholds anywhere, and a candidate should describe the escalation logic rather than reproduce a cut-off.

Note where this places the burden. The system, not the patient, owns the trigger. In ordinary care the patient who is not improving has to come back and say so, which selects for confidence, proximity and money. In a measurement-based system, the review happens on schedule and the absence of improvement declares itself. That single inversion is where most of the equity gain in these models comes from.

Enumerated, not attending	Scheduled, not requested	Caseload, not referral	Obliged, not offered
WHO THE REGISTER COUNTS	WHEN THE REVIEW HAPPENS	HOW THE SPECIALIST SELECTS CASES	WHAT THE MEASURE TRIGGERS

Every tile in that strip names a default that ordinary care sets the other way, which is a useful way to audit an existing service: ask which of the defaults it has actually inverted.

29.4 The joins, which is where the system leaks

Steps are easy to design and joins are hard to run, so the joins are where an examiner will probe. Each transition between steps is a moment at which responsibility can be dropped rather than handed over, and the failure is usually silent because both sides believe the other is holding the patient.

THE JOIN	WHAT MUST BE AGREED BEFORE A PATIENT REACHES IT	HOW IT FAILS IN PRACTICE
Recognition into the programme	Who may enter a patient on the register, and on what basis	Entry depends on individual clinician interest, so the register reflects enthusiasm rather than need
Low-intensity step to structured intervention	The observation that triggers it and who acts on it	The trigger exists but no capacity was commissioned behind it, so the step is nominal
Structured intervention to specialist decision	Whether the specialist reviews the record or must see the patient	The review defaults to an appointment, and the model reverts to ordinary referral
Any step back down, or discharge	Who resumes routine follow-up and with what safety-netting	Discharge is treated as an ending, and relapse re-enters the system as a new presentation

The number of steps is not standardised and varies between programmes; what is invariant is the structure of the join. An answer that describes a named ladder of steps has described one programme. An answer that describes what must be agreed at each transition has described the model.

29.5 The Indian evidence, and what it licenses

The most directly relevant Indian evidence for this construct comes from public primary-care facilities in Goa, where recovery from common mental disorders improved among the patients attending them when collaborative care was delivered by **a trained lay counsellor**. The finding matters for a specific reason that is easy to state and easy to overstate: the care-manager role was filled by a trained lay counsellor rather than by a mental-health professional, in the public sector, in ordinary facilities.

Read carefully, that result licenses a workforce conclusion and not a universal one. Its population is patients with common mental disorders attending public primary-care facilities, so it says nothing about severe mental illness, about facilities without the supervisory structure the trial supplied, or about a counsellor cadre without the training the trial delivered. The delivery agent is the variable of interest, and the supervision is the condition that makes the delivery agent safe. A programme that copies the cadre and drops the supervision has copied the cheap half.

IN INDIA

Where a district cannot place a mental-health professional in every facility, and most cannot, the practical decision this evidence supports is to build the care-manager role from a trained non-specialist cadre with specialist supervision attached, rather than to hold the model in abeyance until specialist posts are sanctioned. The gain was obtained in ordinary public primary-care facilities with a lay counsellor delivering the collaborative intervention, so the design question for a district programme is what training and what supervisory link it can actually sustain, not whether a psychiatrist can be found for every centre.

29.6 What these models are not, and the errors that follow

Collaborative care is not consultation. In consultation the specialist answers a question about a patient and the arrangement ends there. In collaborative care the specialist joins a standing system with a register, a schedule and a named person holding the caseload, and the relationship persists whether or not any particular question is asked. Liaison practice and the taxonomy of liaison service models belong to the Consultation-liaison, geriatric and transcultural psychiatry notes.

Neither model is case management, and the distinction is examinable. Case management and assertive community treatment are service models for a different population, defined by their caseload logic and their target group, and they are taught with their figures in the Schizophrenia: management, antipsychotics and adverse effects notes. Case management appears here only as a component: the care-manager function inside collaborative care borrows the idea of a named person holding a defined caseload, and borrows nothing else.

TRAP **Describing stepped care as "start with the least intensive treatment", full stop.** Candidates reproduce the first half of the definition because it is the memorable half, and lose the marks that sit on the second. The examinable content is the escalation: what is measured, when it is measured, what threshold obliges a change, and who is responsible for making it. A description without the stepping rule is a description of a waiting list.

PITFALL

Commissioning the register and the care manager but not the step above. A programme that can identify and monitor many more people than the specialist tier can absorb generates a queue of correctly identified, correctly measured, non-improving patients with nowhere to go, and the care manager becomes the person who apologises to them. Capacity at the higher step is a precondition for opening the lower one, not an enhancement to be added later.

29.7 Fidelity, and why a model that worked in a trial stops working as a programme

These are complex interventions, and complex interventions degrade in characteristic ways when they scale. A trial supplies training, supervision, protected time and a research infrastructure that measures whether the protocol was followed. A programme supplies none of those unless somebody commissions them, and the components that are dropped first are always the invisible ones. Supervision is invisible. The register is invisible until it is missing. The scheduled review is visible only to the person who did not receive it.

The practical consequence is that a programme evaluation that reports outcomes without reporting **fidelity** cannot interpret its own result. If outcomes are poor and fidelity was poor, the model was not tested. If outcomes are poor and fidelity was good, the model failed in that setting, which is a genuinely different finding with a genuinely different response. Reporting both is not methodological fussiness; it is the difference between fixing the implementation and abandoning the design.

Three components carry most of the fidelity burden in these models and are worth naming as the things to protect when a budget is cut. The first is supervision of the non-specialist deliverer, because it is what converts a trained cadre into a safe one. The second is the scheduled review, because dropping it converts the register into a list. The third is the capacity at the step above, because escalation that has nowhere to go teaches the whole system to stop escalating. A programme that keeps the counselling sessions and loses all three has kept the visible half of the intervention and discarded the mechanism.

29.8 Evaluation, and the quantities the model has to be able to state

These models are unusually easy to evaluate honestly, because their own machinery generates the necessary denominators. The register defines who was in the programme, so proportions can be computed rather than estimated. A service running collaborative or stepped care should be able to state each of the following, and one that cannot is reporting activity.

- How many people were entered on the register in a stated period, and out of what attending population.
- What proportion received a scheduled review at the intended interval, which is the fidelity measure that matters most.
- What proportion of those not improving at review were actually stepped up, which is the single number that distinguishes the model from a menu.
- What proportion of specialist reviews were caseload-based rather than triggered by an individual referral.
- What happened to those discharged, since a model that improves recovery and loses relapse has moved the problem rather than solved it.

Costing and financing sit outside this construct and are taught in the mental health economics and financing material in these notes; the classic economic argument for funding a liaison service, and the observation that Indian cost-effectiveness evidence in this space is scarce, are already made in the Consultation-liaison, geriatric and transcultural psychiatry notes and are not re-derived or re-declared here.

> Collaborative care is defined by a register, a care manager and caseload-based specialist review; stepped care is defined by the rule that forces the step, not by the list of steps.

Special-setting and disaster mental health

30 · Disaster mental health: psychological first aid, phases and intervention models

A disaster does not produce one clinical problem; it produces a population with a changed distribution of need. Most of the people affected will recover without any formal intervention. A minority will develop diagnosable disorder. A separate and often forgotten group already had severe mental illness before the event and has just lost its supply of medicines, its clinic, its

records and its family carers. A group larger than any of these has been displaced, and **displacement** generates its own morbidity from crowding, lost livelihood and lost privacy rather than from the original event. A last group is the responders. Any response designed around only one of those groups fails all the others, and that is the examinable point: disaster mental health is a service-design question about matching a tiered response to a heterogeneous population, not a question about which therapy to give a survivor. The statutory machinery of disaster management, its tiers and its office-holders, and the professional and ethical frame around emergency work are taught in the Forensic ethics, consent and legal procedure notes and are not restated here. The clinical early-trauma intervention itself, with its action principles, its follow-up rule and its standing cautions about early single-session and pharmacological approaches, is taught in the Anxiety, OCD and trauma-related disorders notes. What follows is the population response and the service that has to deliver it.

30.1 Needs assessment as a method, not as an impression

The first deliverable of a mental-health response is not an intervention, it is a description of who needs what. **Rapid needs assessment** is a method with a defined output, and a response that skips it delivers whatever the responding agency happens to be able to provide, which is a common and expensive failure.

A usable assessment establishes several things that cannot be inferred from the event itself. It establishes the size and composition of the affected population, including how many are displaced and where they now are. It establishes what has been lost from the pre-existing service: which facilities are standing, which staff are contactable, which drug supply lines still run, and whether patient records survived. It establishes what the population itself says it needs, which in the acute period is reliably practical rather than psychological. And it establishes what already exists locally, meaning the community structures, faith networks and lay helpers who will still be present long after any external team has left.

The output of that assessment should be a distribution rather than a headline. A single figure for the proportion of a population needing mental-health care is almost always useless in operations, because it merges people who need information and practical help with people who need medication continued and people who need specialist treatment. The three are met by different tiers with different workforces and different costs.

MNEMONIC

Assess, Tier, Train, Track

The sequence a disaster mental-health programme is built in, and the order in which it is usually cut. Assess before responding. Tier the response so each level of need meets the cheapest workforce that can safely meet it. Train and support the people who will still be there. Track what happened, because a response nobody measured teaches the next response nothing.

30.2 The tiered psychosocial response

The organising structure of every credible disaster mental-health response is a **layered response**, in which the widest layer serves everybody affected and each successive layer serves a smaller group with a more specialised workforce. Read from the base upwards, the layers are these.

The base layer is basic services and security delivered in a way that protects dignity: shelter, food, water, safety, and above all information, because uncertainty about missing relatives and about what happens next is one of the largest sources of distress in the acute period. This layer is not usually thought of as mental-health work and is the layer with the greatest population mental-health effect. Above it sits community and family support: restoring contact between separated people, reactivating community structures, supporting normal routines for children, and enabling collective mourning practices. Above that sits focused, non-specialised support delivered by trained and supervised workers who are not mental-health professionals, aimed at those whose distress is not settling. At the apex sits specialised mental-health care for the minority with severe or persisting disorder, and for those whose pre-existing illness has destabilised.

Three properties of the layered structure carry most of the teaching weight. It is not a triage sequence: the layers operate simultaneously, not one after another. It is inverted in cost and in volume, so the layer serving most people is cheapest and the layer serving fewest is dearest, which is exactly why an unplanned response over-invests at the top. And each layer requires a different workforce, so a plan that names the layers without naming who staffs each one has not been costed.

This chapter cannot attribute the layered framework to an issuing body from a source it has opened, and does not do so. The structure is taught here as service reasoning; a candidate who describes the layers and their workforce implications has the substance, and one who attaches a body and a year from recollection has added the one element most likely to be wrong.

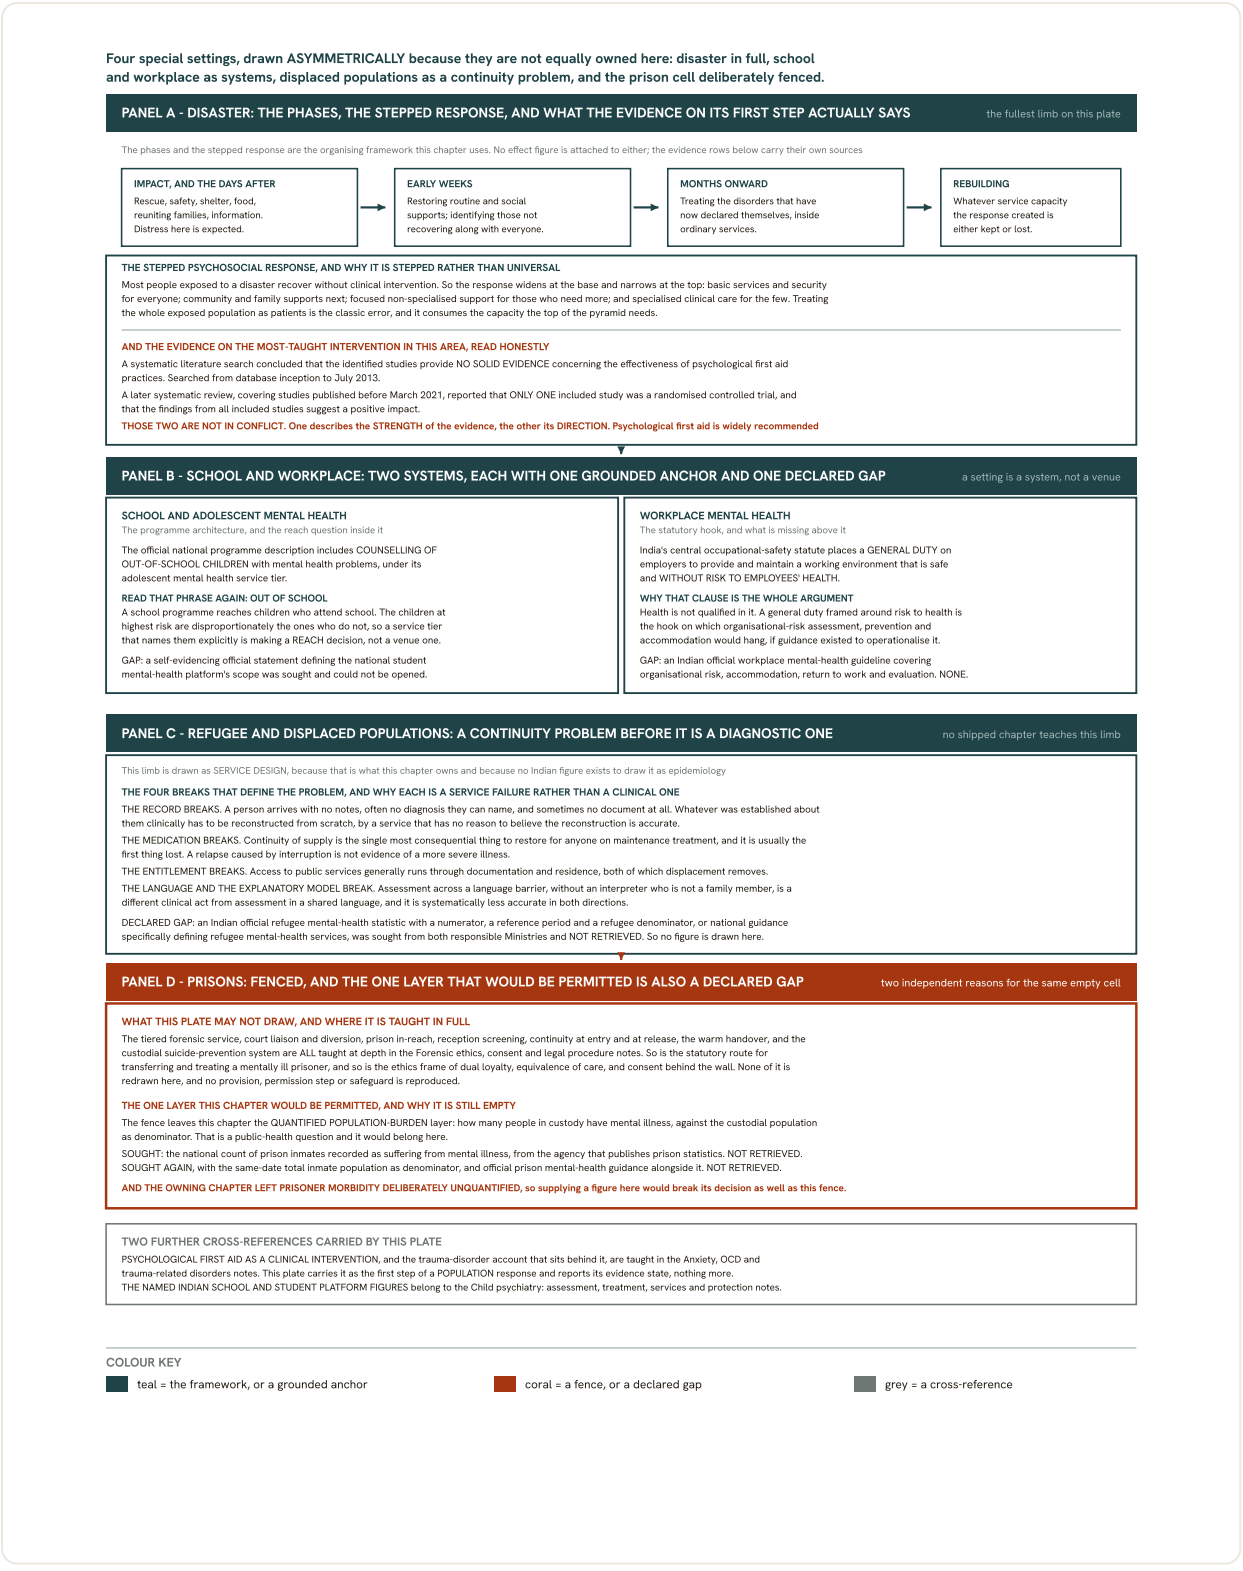


Fig 39.11 Four special settings, drawn to different depths because they are owned to different depths. Disaster carries the phases and the stepped psychosocial response in full, with its first step read honestly: one systematic search found no solid evidence of psychological first aid effectiveness while a later review found only one randomised trial among its included studies and reported that all of them suggest a positive impact, which are statements about strength and about direction rather than a contradiction. School rests on an official service tier that names counselling of out-of-school children, a reach decision rather than a venue one, with the national student platform's scope a declared gap. Workplace rests on a general statutory duty to maintain a working environment without risk to employees' health, with no Indian operational guideline above it. The displaced-population limb is drawn as four breaks in continuity, and no Indian refugee mental-health statistic was retrieved. Prisons are fenced to the Forensic ethics, consent and legal procedure notes, and the one permitted layer is itself a declared gap.

30.3 Where psychological first aid sits, and what it is for

Psychological first aid occupies the focused non-specialised layer. It is a structured, brief, humane response delivered by trained non-specialists to people in acute distress after a critical event, and its purpose is to reduce immediate distress, meet practical needs and connect the person with continuing support. Two boundaries define it more usefully than any description of its content.

The first boundary is that it is not a treatment. It does not target a disorder, it is not indicated by a diagnosis, and completing it is not a therapeutic episode. Its clinical action principles, the rule governing active monitoring and review afterwards, and the standing cautions about what must not be done in the acute phase are taught in the Anxiety, OCD and trauma-related disorders notes, and a candidate should point there rather than reproduce them.

The second boundary is that it is not universal. The people who benefit are those in acute distress, not everyone who was present. A response that pushes an emotional intervention onto every exposed person misallocates the workforce and risks pathologising a normal reaction, and it also crowds out the practical help that the acute period actually demands. The competence being tested when an examiner asks about psychological first aid is usually this discrimination, not the recall of a set of steps.

The delivery implication is the reason the model matters at scale. Because it is deliverable by trained non-specialists, it converts teachers, volunteers, community health workers and relief staff into a usable mental-health workforce for the widest tier of need. That is a workforce multiplication, and it is what makes a population-scale response arithmetically possible at all.

30.4 The evidence base, and how to hold it honestly

This is the part of the topic where candidates most often assert more than the literature supports, and it is worth getting exactly right, because the honest position is more impressive than the confident one.

A systematic search of the literature on this intervention concluded that nothing it identified amounted to solid evidence of effectiveness, which was the basis for its finding that guidelines could not be built on the evidence then available. A later systematic review of the same intervention reported a different but compatible picture: across what it included, randomized evidence was confined to a single trial while the studies as a whole pointed towards benefit.

Recommended widely	**Directionally positive**	**Randomized evidence scarce**	**No solid effectiveness evidence**
STATUS IN DISASTER RESPONSE PRACTICE	SIGNAL ACROSS INCLUDED STUDIES IN THE LATER REVIEW	DESIGN OF THE INCLUDED STUDIES	FINDING OF THE EARLIER SYSTEMATIC SEARCH

Those statements are not in conflict and reading them as a contradiction is the error to avoid. One reports the absence of solid effectiveness evidence; the other reports that the studies that do exist point in a favourable direction while almost none of them are randomized. Both are statements about the evidence base rather than about the intervention, and together they

describe a widely recommended intervention whose recommendation rests substantially on plausibility, consensus and the absence of harm rather than on trial evidence.

RULE **Absence of solid effectiveness evidence is not evidence of ineffectiveness, and it is not permission to assert effectiveness either.** The correct sentence names the intervention, states what the reviews found about the evidence, and gives the reason it is nonetheless recommended: it is low-risk, deliverable at scale by non-specialists, and the alternative on offer in the acute period is frequently nothing at all.

There is a structural reason the evidence is thin, and saying it demonstrates understanding rather than evasion. Randomizing an immediate humane response after a catastrophe is difficult to justify and difficult to execute; the exposed population is displaced and hard to follow; the outcome of interest emerges over a long period; and the intervention is delivered by many different people in many different forms under one name. Those are features of the setting, not failures of the researchers, and they mean the evidence base is unlikely to become deep quickly.

30.5 What changes as the response moves through its phases

A disaster response is not one service running for a long period; it is a sequence of different services with different staffing, and the transitions between them are where continuity is lost. The most important operational fact about the sequence is counter-intuitive: the mental-health workload does not peak when the media attention does. It rises as the acute relief effort demobilises, at precisely the point when external teams withdraw and funding attention moves elsewhere.

PERIOD OF THE RESPONSE	WHAT THE MENTAL-HEALTH SERVICE IS MAINLY DOING	THE CHARACTERISTIC FAILURE
Before any event, as preparedness	Building the plan, the roster and the training so that a response does not have to be invented	No plan exists, so the response is assembled by whoever arrives first
Immediate aftermath	Practical help, information, safety, and restoring supply and records for those already under treatment	Psychological intervention is offered while the practical layer is unmet
Relief and early recovery	Focused non-specialised support, case-finding for persisting distress, and support to responders	External teams deliver in parallel to the local service rather than through it
Later recovery	Specialist treatment for persisting disorder, and handover into the standing service	The external response has already left and the standing service was never strengthened

The service consequence of the last row is the single most useful thing to be able to say about disaster mental health. A response measured by what it delivered during the emergency will look successful and leave nothing behind. A response measured by what the district can do afterwards is a different programme, staffed and budgeted differently from the start.

The first row is the one under a service's own control and the one most often empty. **Preparedness** is the only phase in which decisions can be taken calmly, and it is the phase in which a district can decide, in advance, who leads the mental-health component, which staff are on the roster, what training they will have had, how a drug supply is re-established, where displaced patients will be seen, and what the mutual-aid arrangement with neighbouring districts is. None of those questions gets a good answer in the week after an event, and every one of them determines what the response is able to do.

The transition that matters most is the last one, and it has a name worth using precisely. **Handover** here does not mean transferring a list of patients; it means transferring a functioning arrangement, including who now holds each patient, where the record sits, who supervises the trained non-specialists, and which budget line carries the work from that point. A handover that consists of a summary document is an ending described as a transition.

PITFALL

The parallel service. An external team arrives with its own staff, its own protocol, its own records and its own funding line, works alongside the district service without integrating with it, and departs when the funding ends. Every patient it was holding is then either lost or handed to a service that has never seen them, has no record of the treatment given and was not resourced for the extra load. Working through the existing service is slower at every point and is the only arrangement that leaves capacity behind.

30.6 The responders are an exposed population

The workforce delivering the response is itself part of the affected population, and in a local disaster it is often part of it twice over. Rescue workers, health staff, police, volunteers and administrators are exposed to the event, to its aftermath and to sustained contact with distress, and in a district-level disaster many of them have lost homes and relatives in the same event they are working through.

Responder care is therefore a programme component and not a gesture, and it is designed at the level of the system rather than the individual. The measures that make a difference are structural: rostering that enforces rest rather than relying on individuals to request it; briefing before deployment about what will be seen; a named supervisor with an explicit responsibility for the team's condition; a mechanism to stand somebody down without it being treated as a failure; and access to care that is confidential enough that a local worker will actually use it in a place where everybody knows everybody.

The confidentiality point deserves emphasis because it is where responder programmes most often fail in practice. A support service staffed by the worker's own colleagues, in the worker's own institution, is a service the worker may not use for exactly the problems it exists to address. Designing an access route that sits outside the immediate hierarchy is a service-design decision that costs almost nothing and determines whether the provision is real.

30.7 Special groups the tiered plan has to name explicitly

A layered plan distributes generic need well and handles specific vulnerability badly unless the vulnerable groups are named in the plan itself. Each of the following requires a different arrangement, and each is routinely discovered late.

- People already receiving treatment for severe mental illness, whose immediate need is supply and continuity rather than a new intervention, and whose records may no longer exist.
- Separated and unaccompanied children, for whom reunification and protection are the mental-health intervention and everything else is secondary.
- Older and disabled people, whose access to a distribution point, a shelter or a clinic may be the binding constraint on every other service they need.
- The bereaved, particularly where bodies were not recovered, since the practical and administrative aftermath of an unconfirmed death is itself a prolonged stressor.
- Displaced people living in temporary settlements, where crowding, loss of privacy and loss of livelihood generate new problems that the original event did not.

Naming these groups in the plan does more than allocate resources. It commits somebody to counting them, and a group nobody counts is a group whose needs will be met by whatever is left over.

TRAP **Answering a disaster mental-health question as though the task were to treat post-traumatic stress disorder in survivors.** Candidates go straight to a diagnosis and a therapy, and thereby miss the population structure, the layered response, the people whose established treatment has been interrupted, and the responders. The disorder-level question is a legitimate one and belongs to the chapter that owns acute trauma; the question asked here is what a service does for a population, and the marks are distributed accordingly.

30.8 Evaluation, and what a response should be able to show afterwards

Disaster responses are evaluated less often than almost any other public-health intervention, for understandable reasons and with predictable consequences: each response is improvised, and the improvisation is not recorded, so the next one starts from the same place. Building the evaluation into the plan before the event is the only point at which it is cheap.

What a response should be able to state afterwards is modest and rarely available. How many people were reached at each layer, and out of what estimated affected population. What proportion of those already under treatment for severe mental illness had their treatment continued without interruption, which is the most concrete continuity measure available and the one most likely to reflect real service performance. How many non-specialists were trained, how many were supervised afterwards, and how many were still in post once the response ended. What proportion of responders used the support offered. And what capacity remained in the district when the response closed, compared with what was there before it opened.

That last quantity is the honest measure of a disaster mental-health programme, and it is the one that distinguishes a relief operation from a health-system intervention. An event that leaves a

district with more trained workers, a written plan and a functioning referral route has produced a durable gain. An event that leaves a district with a completed activity report has not.

> The response is layered and simultaneous, psychological first aid is one rung of it delivered by trained non-specialists, and the population that most reliably needs a psychiatrist is the one that already had one before the disaster.

31 · School mental health programmes and adolescent mental health services

The school is the only place where a health system can reach almost an entire age group without anyone having to decide they are unwell. That single property is what makes school mental health a public-health topic rather than a paediatric one, and it is what the marks here are actually about: the school is a delivery platform, and the examinable questions are the ones asked of any platform. Which tier of need is being addressed. Who delivers at each tier. What proportion of the intended population is genuinely reached. Who pays, and for how long. And how anybody would know whether it worked. What a school programme contains, how life-skills teaching is built, how a single session becomes a school culture, and how an identified child is assessed, treated, safeguarded and referred are taught at length in the Child psychiatry: assessment, treatment, services and protection notes, which also carries the Indian delivery platforms with their coverage figures. None of that is reproduced here. What follows is the platform read as a system: tiers, coverage, financing, evaluation, and the adolescents the school platform structurally cannot reach.

31.1 The school as a platform, and what the platform property buys

Every population intervention faces a selection problem. Services delivered in clinics reach the people who present, and presentation is filtered by recognition, permission, money and distance, all of which are unevenly distributed. A **platform intervention** avoids the filter by delivering where the population already is. Schools are the most complete such platform available for young people, and their advantages are structural rather than clinical.

The assembled population is the first advantage: no health-seeking decision is required, so the intervention reaches those who would never have been brought. The second is repeated contact over a long period, which is what any preventive intervention needs and what a clinic almost never has. The third is the presence of a trained adult workforce that already observes the child daily and is capable, with support, of noticing change. The fourth is that the school is itself part of the causal environment, so it can be modified rather than merely used as a venue.

The disadvantages are equally structural and are less often stated. The platform reaches only the enrolled and the attending, so it systematically misses the young people at highest risk. Teachers are not a free workforce and every task added to them displaces another. The school's institutional interest is educational attainment, so a mental-health programme is retained when it is seen to serve that interest and dropped when it is not. And a programme delivered inside a school inherits the school's own hierarchies, so confidentiality is harder to guarantee than in a clinic, which is a matter for the chapter that owns the identified child.

RULE **A school mental-health programme is a platform decision before it is a content decision.** The question that determines its effect is not which curriculum is delivered but which tier of need it is aimed at, who delivers it, and what proportion of the intended age group is actually enrolled, attending and present when it is delivered.

31.2 The prevention tiers applied to a school

The prevention taxonomy is developed as a framework in the prevention material of these notes, and its definitions belong elsewhere again; what belongs here is its application, because the school is the setting in which all its tiers can be seen operating in one building. Applied to a school population, the tiers separate cleanly by whom they are aimed at.

TIER	WHO IT IS AIMED AT IN A SCHOOL	WHO DELIVERS IT, AND WHAT IT COSTS
Universal	Every pupil, irrespective of risk, and the school environment itself	Teachers within the timetable; cheapest per head and dearest in total, because everyone receives it
Selective	Groups with a raised exposure but no identified problem, defined by circumstance rather than by symptom	Trained school staff or a visiting worker; requires a defensible basis for grouping and careful handling of labelling
Indicated	Individual pupils showing early problems that fall short of a disorder	A trained non-specialist with supervision, or a visiting mental-health worker; smallest numbers, highest unit cost
Treatment, which is not prevention	Pupils with an established disorder	The health service, not the school; the school's job here is recognition and a working referral route

The last row is where most school programmes fail, and it is a system failure rather than a content failure. A universal programme in a school raises recognition, and raised recognition generates identified children. If no referral destination exists, the programme has manufactured demand it cannot serve, and the teachers who identified the children carry the consequence. Building the highest tier before opening the widest one is the single most consistent lesson from school programmes in any health system.

MNEMONIC

Enrolled, Attending, Present, Reached

The four filters standing between an intended school population and the population a programme actually serves. Coverage claimed against enrolment is an administrative figure. Coverage claimed against those present when the intervention was delivered is a service figure. The gap between them is where school programme evaluations most often go wrong.

31.3 The Indian delivery platforms, named and located

India does not lack school and adolescent platforms; it has several, they sit in different ministries, and their coverage provisions, training structures and target ages are taught with their figures in the Child psychiatry: assessment, treatment, services and protection notes. This chapter names them so that a reader knows what exists and where the detail lives, and reproduces none of their coverage figures.

- The district mental health programme's school outreach component, which places mental-health work inside schools through the district team and its master-trainer model.
- Rashtriya Kishor Swasthya Karyakram, the national adolescent health programme, which carries adolescent mental health inside a wider adolescent health remit and includes a clinic-based adolescent-friendly service tier.
- The Ayushman Bharat School Health and Wellness Programme, which works through designated teachers delivering scheduled health-promotion sessions within the school timetable.

The system-level observation worth making about that list is the one a candidate can make without any figure at all. The platforms are plural, they sit under different ministries with different reporting lines, and they were designed at different times for different primary purposes. A district therefore does not have one adolescent mental-health programme with a coverage figure; it has several partial provisions whose overlaps and gaps nobody is required to reconcile. Integration across those platforms is the practical Indian problem, and it is an administrative problem rather than a clinical one.

31.4 The adolescents who are not in school

Every advantage of the school platform is conditional on enrolment, and the conditionality is not random. Young people out of school are disproportionately those who are working, married early, disabled, displaced, from the poorest households, or already in difficulty. On any plausible account of risk they are the group with the highest need, and the platform designed to reach adolescents reaches them least.

This is not an oversight in the Indian service description. The **National Mental Health Programme**'s own account of what its district teams deliver on outreach includes the counselling of children who are not in school and who have mental health problems, listed within its adolescent mental health work. The provision is therefore already named at national level, which changes the nature of the question a district faces: it is not whether an out-of-school limb is legitimate, but what contact point it uses.

That contact point is the whole design problem, because there is no assembled population to work with. The available options are all second-best and each carries a different reach and a different cost. Work through the health facility, which returns the selection filter the school platform existed to avoid. Work through the community health worker, which reaches households but adds to an already loaded role. Work through workplaces where adolescents are employed, which reaches the working young but requires an employer's cooperation. Work through

community and religious institutions, which reaches those inside them. Work through a remote or helpline route, which requires a device and a private moment. A district plan that names one of these and resources it has an out-of-school limb; a plan that names none has a school programme and a claim.

IN INDIA

Because the national programme's own outreach description already includes counselling for children outside school who have mental health problems, the practical decision for a district team is not whether to serve this group but which single contact point it will commit to and staff. Choosing one route and running it properly reaches more of this population than distributing the intention across all of them, and it makes the coverage question answerable, because a named route has a denominator that can be counted while a general intention does not.

31.5 Coverage, and the difference between a programme existing and a programme reaching

Coverage is the quantity on which any platform intervention is properly judged, and it is the quantity most often replaced by a proxy. A programme can be sanctioned in a district and absent in its schools. It can be present in a school and delivered to a fraction of classes. It can be delivered to every class and reach only those present on the day. And it can reach everyone present and change nothing, because the tier addressed was not the tier of need.

The honest formulation asks four questions in sequence and refuses to collapse them. What is the intended population, defined by age and by geography rather than by enrolment. What proportion of that population is enrolled and attending, which is an education-system figure the health system must obtain rather than assume. What proportion of those actually received the intervention as designed. And what proportion of those identified as needing the next tier reached it. A programme that can answer only the first two is reporting the education system's performance, not its own.

Districts	Schools	Teachers trained	Adolescents reached
WHAT A SANCTION FIGURE COUNTS	WHAT AN INSTITUTIONAL FIGURE COUNTS	WHAT A TRAINING FIGURE COUNTS	THE ONLY DENOMINATOR THAT ANSWERS COVERAGE

Read that strip from left to right as a descent from what is easy to report towards what a population actually received. Each tile counts a real thing, and only the last one has the population in its denominator. A programme report that stops before the last tile has not answered the coverage question; it has answered an administrative one and left the reader to assume the rest.

TRAP **Quoting the number of schools or districts covered as though it were population coverage.** Districts covered is an administrative reach figure, and its denominator is districts, not children. Schools covered is likewise a figure about institutions. Neither can be converted into a proportion of the adolescent population without the enrolment and attendance data that turn one into the other, and the young people at highest risk sit precisely in the gap between those denominators.

31.6 What an adolescent service tier has to look like to be used

An adolescent service is not a paediatric service with older patients or an adult service with younger ones, and the reasons are about access rather than about diagnosis. Adolescents present late, present through somebody else, and stop attending for reasons that have more to do with the shape of the service than with their symptoms. Designing an **adolescent-friendly service** means removing the specific obstacles that this age group faces, and each of them is a system decision.

Location and hours are the first, because a service open only during school or working hours obliges the young person to explain their absence, which is itself a disclosure. The entry route is the second: a service reached only through a parent excludes the adolescents whose difficulty involves the household, and a service reached only through a psychiatric department attaches a label at the door. Staff attitude is the third and is the one adolescents themselves report as decisive. And a clear statement of what will and will not be shared with a parent or a school is the fourth, since an adolescent who cannot predict where their words will go will manage the risk by saying less. The confidentiality boundary, the limits of the teacher's role and the handling of an identified child are taught in the Child psychiatry: assessment, treatment, services and protection notes and are not restated here.

The last design problem is **transition**. A young person who reaches an adolescent service and remains unwell must eventually move into adult services, and that move is a documented point of disengagement across health systems. Treating transition as an administrative transfer rather than as a planned handover with an overlap period, a named receiving clinician and a shared plan produces a predictable loss of exactly the patients who most needed continuity. A district that builds an adolescent tier without deciding what happens at its upper edge has built a service with an exit and no destination.

31.7 Financing, and why school programmes have a characteristic life cycle

School mental-health programmes fail in a recognisable pattern, and the pattern is financial rather than clinical. A programme is introduced with project or scheme funding that pays for training and for a coordinating post. The training is delivered and counted. The coordinating post ends with the project. The trained teachers are transferred, promoted or retire, and the school reverts. A subsequent evaluation finds no effect and concludes that the intervention does not work, when what happened is that it was never sustained long enough to be tested.

The financing questions that determine whether this happens are worth naming precisely, because they are answerable in advance. Whether the recurring costs, meaning refresher training and supervision, sit on a recurring budget line or a project one. Whether the coordinating role belongs to a post or to a person. Whether the receiving health service at the top tier was funded to take the referrals the programme generates. And whether the education system carries any of the cost, since a programme entirely funded by the health side is a programme the school has no institutional stake in retaining.

PITFALL

Counting teachers trained as the programme's output. Training is an input, it is easy to count, and it is what appears in reports. It tells you nothing about whether anything was delivered afterwards, whether the trained teacher is still in that school, or whether a single child reached a service. A programme reporting only training numbers has reported its expenditure.

31.8 Evaluation, and what this chapter cannot establish

The evaluation of a school mental-health programme should be aimed at the tier it addresses, and mismatched evaluation is the commonest reason a reasonable programme is judged ineffective. A universal intervention should be judged on population-level change and on the school environment, not on the recovery of identified children, of whom it will contain few. An indicated intervention should be judged on the outcome of those identified, and on whether they reached the next tier. Judging a universal programme by clinical outcomes in a handful of identified pupils guarantees a null result whatever the programme did.

Two things this section would ordinarily state are declared as gaps rather than supplied, and both are exactly the kind of fact that feels certain from memory.

The first is the national student mental-health support initiative run from the education side. What was sought was a primary official statement defining its scope for school, college and university students; the official routes approached timed out or returned no retrievable page text, and no such statement was obtained. The construct exists and is named here; its scope, its year and its provisions are not established by anything this chapter has read, and a candidate is better served by saying so than by attributing a remembered description to a ministry.

The second is prevalence. No defensible Indian figure is available here for cyberbullying among school-going children and adolescents, and none is supplied. This matters as method and not only as caution: a prevalence figure measured in one population cannot be transferred to another simply because both consist of school children, and a figure quoted without its survey, its year, its age band and its definition of the behaviour is not a fact about India. The correct answer names the construct, says the Indian figure is not established, and moves to what a service should do about a problem whose size it does not know.

A school programme's effect is set by its tier and its coverage, and its coverage is capped by enrolment, which is precisely why the out-of-school limb is not an afterthought but the part that reaches the highest risk.

32 · Workplace and occupational mental health

Work is where most adults spend most of their waking, non-domestic life, and it is the one setting in which an institution can change a person's exposure by changing its own arrangements. That is why workplace mental health is a public-health topic and not an occupational-medicine footnote: the employer holds levers no clinician holds. It can change how much is demanded, how much control a person has over it, whether support exists, whether the

role is intelligible, and whether the job will still exist next month. A clinician can treat the consequences of all of these and can change none of them. The examinable content sits on that asymmetry, and the commonest failure in an answer is to describe workplace mental health as a set of services offered to employees, which relocates the problem into the worker and leaves the work untouched. What follows is the organisational risk, the prevention structure applied to a workplace, the accommodation and return-to-work system, the manager's role and its limits, and what an Indian programme can and cannot be held to.

32.1 The setting, and the three populations inside it

A workplace, viewed as a health platform, has the same assembled-population advantage as a school and an additional one: the institution has a continuing relationship with the person, a record of their attendance, and a financial interest in their capacity to work. It also has three distinguishable populations, and a programme that does not separate them will be judged against the wrong outcome.

The first population is everybody employed, well and unwell, whose exposure is set by how the work is designed. The second is those with emerging difficulty who are still working, often at reduced effectiveness and usually undisclosed. The third is those absent from work or returning from absence, whose problem is no longer only clinical but also contractual, financial and social. Each requires a different intervention, delivered by a different part of the organisation, and evaluated on a different measure.

Two further groups sit at the edges and are routinely forgotten. Workers in insecure, informal or contract arrangements are exposed to more of the hazards and covered by fewer of the provisions, so a programme built for permanent employees can widen a gap inside its own workforce. And managers, who are asked to carry much of the delivery, are themselves employees with the same exposures.

32.2 The statutory floor, and how far it reaches

India's central occupational-safety statute places a general duty on employers, which is to provide and to keep in place a working environment that is safe and does not put employees' health at risk. The wording of the duty is the load-bearing part, and it is worth reading twice. It concerns the *environment*, not the worker's constitution. It runs to *health*, without a qualifier confining it to physical health. And it obliges the employer to *provide and maintain*, which is a continuing obligation rather than a one-off provision.

A caution about how far that reading may be pressed, and it applies to any statutory duty. The exact strength of a duty, including whatever qualification the statute attaches to it, is the substance of the obligation and not decoration around it; a duty framed absolutely and a duty framed with a practicability qualifier are different claims about what an employer must do. This chapter states the duty and does not reproduce the statutory text, the section locator or the qualifying words, and a candidate arguing a legal point should go to the statute rather than to a summary of it. The disability entitlement frame that sits behind adjustments in employment is taught in the Mental health legislation and forensic psychiatry notes.

RULE **Workplace mental health is a work-design intervention first and a service offering second.** Interventions aimed only at the individual, however good, leave the exposure in place and locate the problem in the person exposed to it. An answer that lists counselling services without addressing how work is organised has described the smaller half of the field.

32.3 Organisational risk, described as properties of work rather than of workers

The **psychosocial hazards** of work are features of how a job is arranged, and they are describable independently of who holds the job. The categories that recur across occupational settings are these, and each has a corresponding organisational remedy that no individual intervention can substitute for.

- Demand: the volume, pace and difficulty of the work relative to the time and resources given for it.
- Control: how much say the worker has over how, when and in what order the work is done.
- Support: whether practical help and supervision are actually available from managers and colleagues.
- Role: whether what the person is responsible for is defined, consistent and free of conflicting instructions.
- Relationships: exposure to bullying, harassment, discrimination and unresolved conflict.
- Change and security: how organisational change is managed, and whether the job itself is precarious.

The important structural claim about these hazards is that they interact rather than add. High demand alongside high control is a demanding job; high demand alongside low control is the arrangement most consistently associated with harm. Effort without corresponding reward, recognition or security behaves similarly. Named theoretical models formalise these interactions and attach evidence to them; this chapter does not name them or their originators, because no claim row here carries them, and a model attributed from memory to the wrong author is exactly the kind of error a viva finds. The categories above are usable without the model names, and an answer that reasons from demand and control to a specific organisational remedy has demonstrated the understanding the model exists to convey.

MNEMONIC

Design, Detect, Adjust, Return

The four functions of a workplace mental-health programme, in the order of their population effect and the reverse of the order in which they are usually built. Design the work so it does less harm. Detect difficulty early and make disclosure survivable. Adjust the job so the person can keep doing it. And build a route back for those who have stopped.

32.4 The prevention tiers applied to a workplace

The prevention taxonomy is developed as a framework in the prevention material of these notes and is applied here rather than restated. Applied to a workplace it separates the interventions by

the level at which they act, and the separation is what makes an evaluation interpretable.

LEVEL AT WHICH IT ACTS	WHAT IT LOOKS LIKE IN A WORKPLACE	WHO HAS TO ACT	HOW IT CHARACTERISTICALLY FAILS
The work itself, before anyone is unwell	Workload and rostering decisions, clarified roles, managed change, an enforced anti-harassment position	Senior management, because these are budget and structure decisions	Delegated to a health or welfare function that cannot change any of them
The whole workforce, as capability	Awareness that reduces the cost of disclosure, and training that makes managers approachable	The organisation, with time released for it	Delivered as a single session and never refreshed, so the culture reverts
Individuals with emerging difficulty	Early conversation, confidential access to help, adjustment before absence occurs	Line managers, supported by an occupational health or counselling route	Access exists but disclosure is unsafe, so nobody uses it until they are absent
Those already absent or impaired	Structured return to work, graded duties, sustained adjustment, continuity with treatment	Manager, occupational health and the treating clinician together	Handled as an attendance problem by a process designed for discipline

The job	**The person**	**Disclosure cost**	**Return route**
WHAT AN ORGANISATIONAL INTERVENTION CHANGES	WHAT AN INDIVIDUAL INTERVENTION CHANGES	WHAT DETERMINES WHETHER EITHER IS USED	WHAT DECIDES IF ABSENCE BECOMES EXIT

32.5 Adjustment, and why it is cheaper than everything downstream

Adjustment means changing the job so that a person with a mental-health difficulty can continue to do it, and its most useful property is that it is usually inexpensive and reversible. The adjustments that matter in mental health are unglamorous and mostly concern time, predictability and load: altered start and finish times so that treatment appointments or medication effects are accommodated; a temporary reduction in the most demanding component of the role rather than in the role as a whole; a quieter or more predictable working location; a phased increase in duties; a named person to check in; and a temporary change in the frequency of supervision so that difficulty is caught before it accumulates.

Two design principles carry most of the effect. The first is that adjustment should be available before absence occurs, not only as a condition of return, because the cheapest point to intervene is while the person is still at work. The second is that an adjustment must have a review date and a person who owns it, or it either lapses silently or becomes permanent by default, and both outcomes damage the arrangement's credibility for the next person who needs one.

IN INDIA

Because the employer's general duty concerns a working environment that is safe and without risk to employee health, and that duty is not written as a duty about physical health alone, the practical decision for an Indian organisation is to locate mental-health provision inside the occupational safety and health management system it already runs rather than beside it as a welfare activity. That placement is not cosmetic: it puts the provision where the statutory duty, the reporting line and the standing safety committee already sit, which is what converts an initiative into an obligation somebody is answerable for.

32.6 Return to work, treated as a system rather than a date

Return to work is where workplace mental health becomes measurable and where most organisations perform worst, because absence is managed by processes designed for a different purpose. The default machinery is **attendance management**, whose logic is verification and escalation, and applying it to a mental-health absence reliably delays return and sometimes prevents it.

A functioning return system has a small number of features and they are all about contact and structure. Contact during absence is maintained by a named person on an agreed basis, since silence is read as abandonment and repeated formal letters are read as threat. A return is planned before it happens, with the returning person present at the planning. Duties are **graded** rather than restored in full on the first day. The adjustments are written down, given a review point, and communicated to whoever needs to know and to nobody else. And the return is reviewed after a period rather than treated as complete on arrival, because relapse in the first weeks back is common and is best met with a further adjustment rather than a further absence.

PITFALL

Requiring full fitness before return. The demand looks prudent and is counterproductive: it converts a person who could do part of the job now into a person doing none of it indefinitely, removes the structure and contact that support recovery, and lengthens absence to the point at which returning at all becomes unlikely. A **capacity-based return**, which asks what the person can do and adjusts around what they cannot, is both safer and faster.

32.7 The manager's role, its limits, and the confidentiality problem

Most workplace mental-health programmes rest on line managers, and most manager training fails by being too ambitious about what a manager is for. A manager is not being trained to assess, diagnose, counsel or treat, and training that implies otherwise produces either an anxious manager who avoids the conversation or a confident one who does something inappropriate.

The manager's job has three parts and stops there. Notice a change in the person's work or presentation. Have a conversation that makes disclosure safe and is about work rather than about diagnosis. And connect the person to a route, whether that is occupational health, a counselling service or their own doctor, while adjusting what can be adjusted meanwhile. Training aimed at

those three does measurably more than training aimed at recognition of disorders, and it is also the version a manager can practise without exceeding their competence.

Confidentiality is the structural problem that determines whether any of this is used. An employee discloses to someone who also decides their appraisal, their duties and possibly their continued employment, and the same information that unlocks an adjustment can be read as evidence of unreliability. The resulting **cost of disclosure**, meaning what the employee reasonably expects to lose by speaking, is the single variable that determines whether any provision in the organisation is used at all. The design answer is to separate what the manager needs to know from what they must not: a manager needs to know what adjustment is required and for how long, and does not need a diagnosis, a treatment or a prognosis. Where an **occupational health function** exists, its correct output to a manager is a statement of capacity and adjustment, not a clinical account. Where a service is delivered by the employer's own staff, the same problem arises inside the service, which is why an access route independent of the reporting line is a design requirement rather than a courtesy.

TRAP **Reporting uptake of an employee counselling service as the outcome of a workplace mental-health programme.** Uptake is an activity figure with no population denominator attached, and it can rise for good reasons or bad ones: a workforce under increasing strain generates more uptake, as does a service that has become easier to reach. It says nothing about exposure, about absence, about return, or about the people who did not disclose. The outcomes that answer the question are absence and return-to-work measures, the proportion of adjustments requested that were granted, and any measure of the psychosocial hazards themselves.

32.8 Evaluation, and what is not established for India

The evaluation problem here is that the easily available measures belong to human resources and the meaningful ones do not exist unless somebody builds them. Absence data are collected everywhere and are a poor single measure, because absence falls both when people get better and when they become too insecure to take leave. A defensible evaluation therefore reads several quantities together: absence and its duration, the proportion of absences that end in return rather than exit, the proportion of requested adjustments granted and sustained to their review point, and a repeated measure of the psychosocial hazards themselves, which is the only quantity that reports on the intervention that matters most.

Two things a section on Indian workplace mental health would ordinarily supply are declared here as gaps. There is no primary-official Indian workplace mental-health guideline available to this chapter addressing organisational risk, prevention, reasonable accommodation, return to work, manager training and programme evaluation. What was sought was exactly such an instrument from the labour and health ministries and their technical bodies; the official sources approached either carried only a bare mention of workplace stress management without any framework, or established a general health duty without an express mental-health framework, or failed to open. The construct is named here without a national instrument behind it, and a candidate who says that plainly is answering correctly.

There is likewise no national coverage figure available: how many Indian workplaces run any mental-health provision, what proportion of the workforce it reaches, and what proportion of that workforce is in the informal arrangements that no such programme touches are all unquantified here. The absence is itself instructive, because a workforce whose majority is informally employed cannot be covered by an employer-based model however good that model is, and a national workplace mental-health strategy that addresses only formal employment addresses a minority of the working population.

> The employer's duty runs to the working environment and to health without qualification, so the first intervention in workplace mental health is always to the design of the work, and every service offered to the worker is downstream of it.

33 · Prison, refugee and other special-population mental health

A special population, in service terms, is not a population with an unusual diagnosis; it is a population whose access to care is controlled by something other than its need. For prisoners that something is custody. For refugees and displaced people it is legal status, documentation and mobility. For others it is the absence of an address, a guardian or an entitlement. The clinical material is largely the same material taught everywhere else in psychiatry; what changes, and what carries the marks, is that the ordinary route to care has been interrupted by an institution or a legal category, so the service has to be redesigned around the interruption. This section teaches the displaced-population limb in full, and for prisons only the population and system questions: what a prison mental-health figure would have to be to mean anything, and how prison mental-health systems differ across countries. Custodial law, prisoner rights, the ethics of practice behind the wall, the design of an in-reach service, continuity at entry and at release, and the custodial suicide-prevention system are taught at depth in the Forensic ethics, consent and legal procedure notes, and none of them is restated here.

33.1 What makes a population special, stated as a rule about access

The useful generalisation is that in each of these populations a third party stands between the person and the service, and that third party's interests are not identical to the patient's. A prison authority is accountable for security. An immigration authority is accountable for status determination. A shelter or an institution is accountable to whoever funds it. None of them is primarily accountable for the person's mental health, and each can, without malice, produce a decision that removes the person from care.

Three consequences follow and they recur across every population in this section. Continuity is fragile, because the person can be moved, released, transferred or deported without reference to their treatment. Consent and confidentiality operate under constraint, because the person's circumstances limit what a refusal can mean and who else has an interest in the information. And ordinary measures of service performance mislead, because the denominator is a population the health service does not control and often cannot count.

RULE **Where a third party controls access, service design is mostly about the joins with that third party and only secondarily about the clinical content.** The question that determines whether a special-population service works is not what treatment it offers but what happens to a patient at the moment the custodian, the authority or the institution makes a decision about them.

33.2 Displacement, read as a sequence of exposures rather than an event

Refugee and displaced-population mental health is often taught as though the trauma were the flight. That framing is incomplete and, for service design, misleading. The exposures are distributed across a sequence, and the parts a receiving service can actually change lie at the end of it.

Before departure sit the conditions that produced the displacement: conflict, persecution, disaster, destitution, and the losses accumulated while deciding to leave. During movement sit the hazards of the journey itself, separation from family, and, for many, exploitation and violence en route. After arrival sit a different set of exposures that are chronic rather than acute: uncertain legal status, prohibition or absence of work, dependence, crowded or temporary accommodation, discrimination, and the sustained absence of information about what will happen next.

The post-arrival exposures are the ones a health service and a policy can modify, and they behave differently from the pre-arrival ones. The events before and during flight are fixed and historical. The conditions after arrival are ongoing and administrative, and prolonged uncertainty about status functions as a continuing stressor rather than a background condition. That distinction is why an intervention aimed only at past trauma, in a person whose present circumstances remain unresolved, so often disappoints: it treats the part of the sequence that has ended and leaves the part that is still happening.

Two further groups belong in this limb and are frequently excluded from it. **Internally displaced people**, who have not crossed a border, are usually more numerous than refugees and are covered by fewer specific provisions. And people whose displacement is old, meaning long-settled populations in **protracted situations**, have a different problem again: their acute needs have passed and their structural disadvantage has become permanent.

33.3 What a service for displaced people actually has to solve

The design problems are concrete and are mostly not clinical. Each of them has a service answer, and a plan that has not answered them will deliver care to whoever happens to arrive able to use it.

- Entitlement: what a person is allowed to receive depends on a legal status that may be undetermined, contested or absent, so the service must decide in advance what it will provide to somebody whose status is unresolved.
- Language: assessment through an untrained **interpreter**, or through a family member, changes what is said and what is recorded, and interpreting capacity is the single most common binding constraint on the quality of a refugee mental-health assessment.

- Documentation: people arrive without records, so treatment history, previous diagnoses and current medication have to be reconstructed, and continuing supply of an existing medicine is frequently the most urgent clinical task.
- Mobility: populations move between camps, cities and countries, so a plan requiring several appointments at one location may be unusable, and a **portable record** held by the patient is worth more than a good record held by the clinic.
- Calibration: distress is expected and is not by itself a disorder, so a service must be able to distinguish the person who needs housing, information and a resolved status from the person who needs treatment, without dismissing the second or medicalising the first.

MNEMONIC

Status, Speech, Supply, Setting

The four things to establish before any assessment of a displaced person is worth recording. What legal status they hold and what it entitles them to. In what language the encounter will actually happen, and through whom. What treatment was already running and whether it can be continued. And what their living conditions now are, since those are the exposure a service can still change.

PITFALL

Building a separate refugee mental-health clinic funded by a project. It is visible, it is fundable, and it produces a parallel service with its own records, staff and closing date, attached to a population that is mobile and whose entitlement to the mainstream service was never resolved. When the project ends, the patients belong to nobody. Strengthening the general service to receive this population is slower, less visible and the only version that survives.

33.4 What is not established about refugee mental health in India

This is a section where a confident figure would be easy to produce and impossible to defend, so the position is stated plainly. No Indian primary-official refugee mental-health statistic is available to this chapter, meaning no figure carrying a numerator, a reference period and a refugee denominator; and no national guidance specifically defining refugee mental-health services was retrieved either. What was sought was exactly that, from the home and health ministries. The official sources reached described camps, relief schemes and legal procedure, and contained no mental-health statistic and no mental-health-specific guidance for this population.

The absence is worth teaching as content rather than passed over, because it explains a real feature of the field. India hosts displaced populations of several origins under several different administrative arrangements, and a person's entitlement can depend on which arrangement they fall under rather than on what they need. Where no single national mental-health provision defines the service, what a displaced person receives is determined locally, by whichever combination of the general health service, a state arrangement and a non-governmental provider happens to exist where they are. A candidate who can say that, and can say that the national figures do not exist, is describing the system accurately. A candidate who quotes a prevalence for

refugee mental disorder in India is quoting something no Indian official source in this chapter's reach has published.

Not established	Not retrieved	Not retrieved	Not available
INDIAN REFUGEE MENTAL-HEALTH STATISTIC WITH A REFUGEE DENOMINATOR	NATIONAL GUIDANCE DEFINING REFUGEE MENTAL-HEALTH SERVICES	NATIONAL COUNT OF INMATES RECORDED AS MENTALLY ILL	SAME-DATE INMATE POPULATION AS A DENOMINATOR

33.5 Prisons, and the boundary of what belongs here

Prison mental health is taught as a service in the Forensic ethics, consent and legal procedure notes, and the boundary is worth stating so a reader does not go looking for it here. That chapter carries the custodial service architecture and the standards it is held to. The statutory route by which a mentally ill prisoner is transferred and treated is named here only as the dependency it is, since a service whose transfer route does not work has no upper tier, and the reader is sent to the Forensic ethics, consent and legal procedure notes for the provision itself. The rights and ethics frame governing care behind the wall is likewise named here only as the standard against which any prison mental-health system is measured, and is taught in the same place.

What remains for a public-health chapter is narrower and genuinely distinct: the question of what a prison mental-health figure would have to be to mean anything, and the comparison of prison mental-health systems across countries. Both are population and system questions rather than service-design questions, and neither is taught elsewhere in these notes.

33.6 The numbers problem, and what an honest prison figure would require

It is well established, and stated in the chapter that owns the custodial service, that psychiatric morbidity in custody runs far above the community rate. Nothing in this section contradicts that direction. What this section does is explain why attaching a number to it is harder than it looks, and that explanation is examinable in its own right.

Start with what is routinely published. Prison administrations publish **administrative counts**, and a count of inmates recorded as mentally ill is a count of recognition and recording, not a count of illness. It depends on who did the recognising, with what training, at what point in the sentence, using what threshold, and on whether recording it had any consequence for the institution. A system with no reception assessment records few cases; a system with an active assessment records many; and the difference between the numbers describes the two systems, not the two populations.

Then note the denominator problem, which is specific and fatal to casual comparison. A prison population is a **flow** as much as a **stock**: the number held on a given date is not the number who passed through in a year, and the two differ most for the remand population, which is also the population with the highest turnover and often the highest morbidity. A numerator collected across a period, divided by a population counted on one date, produces a figure that is not a proportion of anything.

An Indian figure that would satisfy this chapter's standard therefore needs a numerator with a stated case definition and ascertainment method, a same-date or same-period denominator from the same source, and a stated reference period. That figure was sought here from the national prison statistics published under the home ministry by the National Crime Records Bureau, together with the corresponding total inmate population. It was not obtained: the national publication timed out on repeated attempts, and where an official route did serve, it exceeded the retrieval limit so that no verbatim table and row could be quoted. A related route supplied a total inmate population without the mental-illness numerator that would have to sit above it. No figure is printed here, and the correct answer in a viva is that the direction is established, the Indian proportion is not, and the reason is a numerator and denominator problem rather than an absence of interest.

TRAP **Quoting an administrative count of prisoners recorded as mentally ill as though it were a prevalence.** The count is a recognition figure produced by the prison system's own procedures, its denominator is usually a single-date headcount while the numerator accrues over a period, and the case definition is rarely stated. Candidates reach for it because it is the only number available. The stronger answer names the count for what it is, states that morbidity in custody runs far above the community rate, and explains what would have to be true for a real proportion to be computed.

33.7 How prison mental-health systems differ across countries

Comparing systems rather than services is the other thing this section can legitimately do, and the comparison is more instructive than a table of prevalences would be, because the structural choices explain most of the variation in what is delivered.

AXIS ON WHICH SYSTEMS DIFFER	ONE ARRANGEMENT	THE OTHER ARRANGEMENT	WHAT THE CHOICE DETERMINES
Who is accountable for health inside the walls	The prison administration provides health care itself	The health ministry provides it, inside a facility run by another department	Whether clinical standards and staffing follow health-service norms or custodial ones
Who employs the clinicians	Employed by the custodial service	Employed by an external health provider working in the prison	Independence of clinical judgement, and how easily posts can be recruited to
Where the sickest are treated	Within dedicated units inside the custodial estate	By transfer out to health-service facilities	Capacity limits, waiting for transfer, and who carries the risk while waiting
How coverage is reported	Activity counts from the custodial system	Health-service reporting with a defined population	Whether coverage can be compared with the general population at all

The last axis is the one that makes cross-national comparison so difficult, and it is worth saying explicitly. Where prison health is reported through the custodial system, its figures are not

commensurable with the health system's, so international differences in reported prison morbidity partly measure differences in reporting architecture. A comparison that does not say which architecture each country uses has compared reporting systems and called it epidemiology.

33.8 The other populations, and the structure they share

The bank of special populations is longer than these two, and the shared structure introduced at the start is what makes them teachable together rather than as a list. In each case a third party or a legal category stands between the person and ordinary care, and the service response is a set of joins rather than a new treatment.

People without a home have no address, which is the hidden requirement in most systems for registration, appointment letters, follow-up and medicine supply. People in institutions of any kind have their access mediated by the institution's staff, so their care depends on somebody else recognising a need and acting on it. People who have been trafficked or are in exploitative work have their movement controlled by another person, so confidentiality and the ability to attend are both compromised. Migrant workers within a country lose entitlements attached to their place of origin without gaining those attached to their place of work. In every case, the correct question is the same: at which administrative moment does this person fall out of care, and what would have to exist for them not to.

That question generalises, and it is the reason this material belongs in a public-health chapter at all. Special-population psychiatry is not a set of exotic clinical entities; it is the ordinary service, examined at the points where its assumptions about address, identity, entitlement, language and freedom of movement quietly stop being true.

> In every special population the third party who controls access is the design problem, and the figures these populations generate describe the recognition system that produced them before they describe the population itself.

Prevention, economics and digital mental health

34 · Levels of prevention (primary, secondary, tertiary) applied to mental health

A prevention taxonomy is a routing device, and it earns its place only because it stops the same argument being made several times over. Community psychiatry is full of interventions that sound alike and act at completely different points: a public awareness campaign, a training programme for general doctors, a helpline, a restriction on access to a lethal method, a supported employment service. Placed on the prevention axes, each one declares whom it acts on and at what point in the course of illness it acts, and once that is declared the argument about it becomes tractable. That is the work this section does for the rest of the chapter: it states the mapping onto mental health and then routes each application to the topic that owns it, so that no application is explained twice. The definitions of the terms themselves, with their lineage and the caution against reading the axes as one, are taught in the Intellectual disability and autism

spectrum disorder notes. The machinery of screening, meaning the criteria a screening programme must satisfy and the biases that inflate its apparent benefit, belongs to the Research Methods, Biostatistics and Epidemiology notes. Neither is restated here.

34.1 What the taxonomy is for, and the two questions it asks

The prevention vocabulary exists because two entirely different questions can be asked of any preventive intervention, and answering only one of them produces confusion that looks like disagreement.

The first question is temporal: at what point in the natural history of the disorder does this intervention act? Before it exists, once it exists but before it has declared itself, or after it is established and the target is now disability rather than onset? That axis yields primary, secondary and tertiary prevention.

The second question is about the recipient: who is offered this intervention, and on what basis were they selected? Everybody in a population, a subgroup at raised risk, or an individual already showing early signs? That axis yields **universal**, **selective** and **indicated** prevention.

The axes are independent, and that independence is the entire reason both exist. An intervention has a position on each, and knowing one position tells you nothing reliable about the other. Reading them as a single ladder is a well-known error with its own exam trap, and both the definitions and that trap are taught in the Intellectual disability and autism spectrum disorder notes. What this section supplies is the mental-health mapping and the routing that follows from it.

MNEMONIC

Whom, When, What, Whether

The four questions that make a prevention claim assessable. Whom does it act on, and how were they selected. When in the course of the disorder does it act. What does it actually consist of, delivered by whom. And whether anybody measured the outcome it is supposed to change, in the population it was aimed at.

34.2 The temporal axis, mapped onto mental health

Mapping the temporal axis onto psychiatry is less tidy than mapping it onto infectious disease, and the untidiness is itself worth understanding rather than glossing.

Primary prevention in mental health acts on determinants and exposures before any disorder exists. In practice this means the material conditions and the social environment, most of which sit outside the health sector: household poverty, violence in the home, adverse experiences in childhood, alcohol availability, working conditions, and access to lethal means. The awkward feature is that almost none of these is a psychiatric intervention, and almost all of the effective levers belong to other ministries. That is not a reason to demote primary prevention; it is a reason to say plainly that a psychiatric service's primary-prevention role is largely advocacy, intersectoral work and the design of population-level restrictions, not clinical activity.

Secondary prevention acts once the disorder exists but before it has declared itself or done its damage, and its content in mental health is early detection followed by prompt effective treatment. This is where most of a health service's preventive effort actually sits, because the interval between onset and treatment in psychiatry is long and the harm accrues during it. Note the boundary carefully: the machinery of screening as a discipline, the criteria a screening programme must meet and the biases that make an unevaluated programme look successful, are taught in the Research Methods, Biostatistics and Epidemiology notes. The stage-by-stage design of a screening-to-treatment pathway, and the rule that a screen without a named destination is not a service, are worked end to end in the Perinatal/women's mental health and emergency psychiatry notes. What belongs here is only the placement of secondary prevention as a tier, and the mental-health examples of what is looked for and by whom.

Tertiary prevention acts after the disorder is established, and its target is disability, relapse and social exclusion rather than onset. In psychiatry this is the largest of the three limbs by volume of work, because the conditions that dominate a service are long-term and their principal harm is functional. Almost everything a rehabilitation service does sits here, and it is rarely called prevention, which is precisely why it loses arguments about resources to interventions that carry the word.

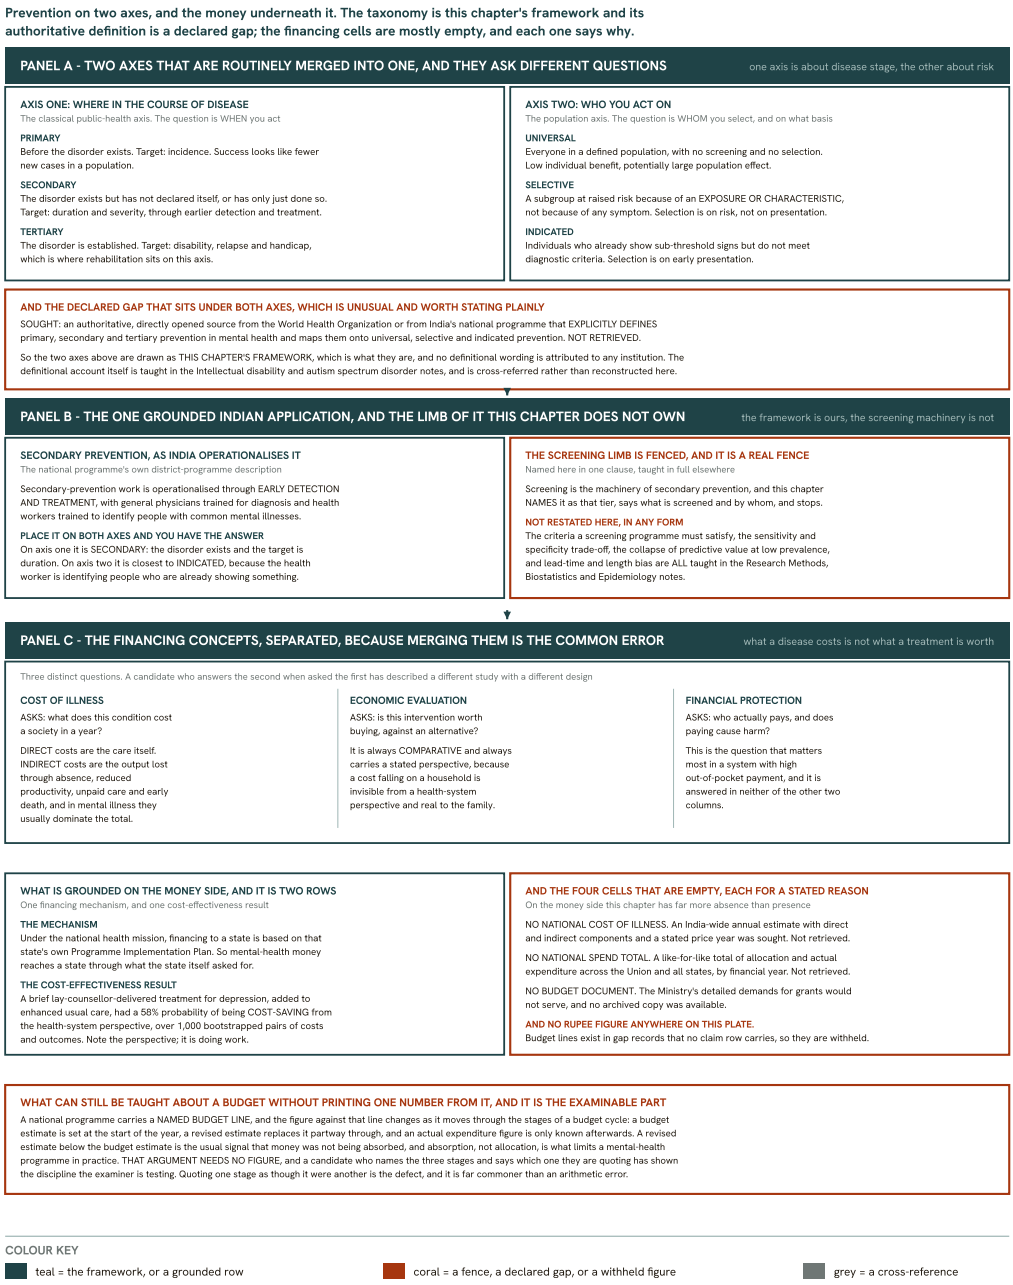


Fig 39.12 **Prevention on two axes, and the financing underneath, with the empty cells labelled.** The two axes ask different questions, disease stage against risk basis, and they are drawn as this chapter's framework because an authoritative source explicitly defining primary, secondary and tertiary prevention in mental health and mapping them onto universal, selective and indicated prevention was sought and not retrieved. The one grounded Indian application is secondary prevention operationalised as early detection and treatment through trained general physicians and identifying health workers; the screening machinery under it is fenced to the Research Methods, Biostatistics and Epidemiology notes. On financing, state money flows through the state's own Programme Implementation Plan, and a brief lay-counsellor-delivered treatment had a 58 per cent probability of being cost-saving from the health-system perspective across 1,000 bootstrapped pairs. No rupee figure appears: no national cost-of-illness estimate and no like-for-like national spend total were retrieved, and the budget lines that exist sit in gap records no claim row carries, so the budget argument is taught through its stages instead.

34.3 The recipient axis, and what it changes about delivery

The second axis is about selection, and its practical value is that it predicts the delivery problem before the intervention is designed. A universal intervention has no selection step at all, so it needs no case-finding apparatus and no consent to being identified as at risk, but it must be cheap per head because everybody receives it and it must be acceptable to people who will not benefit. A selective intervention requires a defensible basis for defining the subgroup, which means either a routinely available marker or a screening step, and it carries a labelling cost that the universal version does not. An indicated intervention requires individual identification, which is the most expensive selection of all and generates the largest ethical obligations, because a person identified as showing early signs is now known to the service and must be offered something.

The examinable consequence is that the cost profile inverts along this axis. Universal interventions have the lowest unit cost and the highest total cost. Indicated interventions have the highest unit cost and reach the fewest people, but concentrate the benefit where risk is highest. Neither is superior in principle, and a programme that has chosen one without stating why has usually chosen by convenience.

Whole population	**Raised-risk subgroup**	**Early signs, identified**	**Established disorder**
WHO A UNIVERSAL INTERVENTION IS OFFERED TO	WHO A SELECTIVE INTERVENTION IS OFFERED TO	WHO AN INDICATED INTERVENTION IS OFFERED TO	WHO SITS OUTSIDE PREVENTION ALTOGETHER

The fourth tile is there deliberately. People with an established disorder are not the recipients of prevention on the second axis at all; they are the recipients of treatment, and of tertiary prevention on the first axis. Keeping that tile in view is what stops a service describing its ordinary clinical work as prevention because the word attracts funding.

34.4 The routing table, which is what this section is for

Every substantial preventive activity in community psychiatry is taught somewhere, and the taxonomy's job is to say where. The table below is the chapter's routing device: it places each application on the axes and names where it is developed, so that the argument is made once, in full, in the right place.

APPLICATION	TEMPORAL TIER	RECIPIENT AXIS	WHERE IT IS DEVELOPED
Restricting access to lethal means	Primary	Universal	The means-restriction and gatekeeper material of these notes, with the population argument in the Somatic-symptom, sexual, impulse-control disorders and suicide notes
Responsible media practice on suicide	Primary	Universal	The means-restriction and gatekeeper material of these notes

APPLICATION	TEMPORAL TIER	RECIPIENT AXIS	WHERE IT IS DEVELOPED
Gatekeeper training for non-clinicians	Secondary	Selective, by role rather than by risk	The means-restriction and gatekeeper material of these notes
A national suicide prevention strategy as a whole	Spans all three	Spans all three	The national suicide prevention material of these notes; the three-tier grid worked through suicide exemplars is in the Somatic-symptom, sexual, impulse-control disorders and suicide notes
Anti-stigma work aimed at the general public	Primary, acting on a determinant of access	Universal	The anti-stigma intervention material of these notes
Mental health literacy and awareness campaigns	Primary and secondary, since recognition shortens delay	Universal	The mental health literacy material of these notes
School-based programmes	Primary, with an indicated limb inside them	All three, in one setting	The school and adolescent material of these notes; the Indian platforms are in the Child psychiatry: assessment, treatment, services and protection notes
Workplace programmes	Primary through tertiary, depending on the component	All three	The workplace material of these notes
Early detection by trained general doctors	Secondary	Indicated, on presentation	This section, and the primary-care integration material of these notes
Screening pathway design in one population	Secondary	Universal or selective, by design	The Perinatal/women's mental health and emergency psychiatry notes
The methodology of screening and its biases	Applies to secondary	Not applicable	The Research Methods, Biostatistics and Epidemiology notes
Rehabilitation, supported employment and community services	Tertiary	Established disorder, so outside the recipient axis	The rehabilitation and community-care material of these notes

RULE **Placing an intervention on the two axes tells you what evidence it needs and what denominator it must report against; it tells you nothing about whether it works.** A tier is a description, not a recommendation. Every row above still has to be argued on its own evidence in the place where it is developed, and the taxonomy's only claim is that the argument belongs in one place.

34.5 How the Indian programme operationalises the secondary tier

The abstraction becomes concrete at exactly one point in the Indian system, and it is worth stating precisely because it is the answer to the most likely viva question on this topic. In **India's National Mental Health Programme**, the secondary tier is operationalised as detecting illness early and treating it, with general physicians trained to diagnose and treat and health workers trained to identify.

Read the structure of that arrangement, because it explains the whole design. The Indian secondary tier is delivered not by an instrument but by a workforce: the detection step is a trained general doctor or health worker recognising a presentation in the course of ordinary contact, and the treatment step follows immediately from the same training. There is no separate screening programme with an invited population and a defined interval; there is an opportunistic detection system riding on existing clinical contact. That is a defensible design in a system with limited specialist capacity, and it has a specific consequence: the capacity of the secondary tier is set by training throughput and by supervision, not by procurement.

IN INDIA

Because the national programme puts its secondary tier in the hands of trained general physicians and identifying health workers rather than in a screening apparatus, a district that wants to strengthen this tier should be counting how many of its general doctors and frontline workers have actually been trained, how recently, and whether anyone supervises them afterwards. Buying or adopting an identification tool changes nothing if the person holding it has no training and no route for a positive finding. Training throughput and supervision are the levers this design exposes, and they are the ones a district can actually pull.

PITFALL

Treating the training as complete once it has been delivered. Detection skill in non-specialists decays without reinforcement, and the trained doctor is transferred, promoted or retires while the programme's records continue to count them. A district reporting cumulative numbers trained is reporting a historical total, not a present capacity, and the gap between those two quantities widens every year the programme runs.

34.6 Tertiary prevention, and why it loses the argument it should win

The largest preventable burden in psychiatry is disability in people who already have a diagnosis, and almost nobody calls the work that reduces it prevention. Relapse prevention, continuity of medication, family support, supported housing, supported employment and the

maintenance of social role are all tertiary prevention, and each prevents a specific, countable harm.

- A readmission that did not happen, which is the outcome most services already record and least often claim.
- A job lost during a relapse, and with it the income and the routine that supported recovery.
- A tenancy or a place in the family home lost while somebody was too unwell to hold it.
- A relationship that ended because nobody supported the family through an episode.
- A contact with the criminal justice system that followed from untreated illness rather than from anything else.

Each item on that list is countable, is already recorded somewhere, and has a cost attached to it in some department's accounts.

The reason this limb is systematically under-argued is partly linguistic and partly budgetary. Linguistically, prevention is heard as meaning stopping something from starting, so activity in people who are already ill is coded as treatment and competes for treatment budgets rather than for prevention ones. Budgetarily, the benefits appear as costs that did not occur, in a budget line different from the one that paid for the service, and often in a different department altogether. A supported housing placement that prevents several admissions saves money in the hospital budget and spends it in the community one, so the service that produced the saving does not receive it.

The counter-argument a candidate should be able to make is that the tertiary tier is the only one whose effect can be demonstrated in a defined denominator over a short period, because the population is already enumerated and the outcomes are already recorded. Where a primary-prevention claim requires a large population and a long follow-up, a tertiary-prevention claim can be tested on a service's own caseload against its own readmission record. The economics of that argument, and the methods for making it, are developed in the mental health economics and financing material of these notes.

34.7 What this chapter cannot establish, and how to answer without it

One thing an authoritative section on this topic would ordinarily supply is declared here as a gap. No authoritative source could be opened that explicitly defines primary, secondary and tertiary prevention in mental health and maps them onto the universal, selective and indicated scheme in a single statement of the taxonomy. What was sought was exactly that: an international agency publication on the prevention of mental disorders, several of its repository records, secondary reproductions of the taxonomy in the indexed literature, and Indian operational guidance for mental health at the primary-care level. The agency records returned access errors, the literature routes returned verification pages rather than article text, and the Indian operational documents supplied promotion, detection, screening and referral examples without labelling the full taxonomy.

The honest consequence is stated rather than concealed: the taxonomy is taught here as the standard framework it is, its definitions are cross-referred to the chapter that carries them, and no issuing body or year is attached to it in this section. A candidate asked to attribute the framework should say which chapter of their own reading carries the definitions and should not attach an agency and a date from memory, because a framework misattributed to a body that did not issue it is a discrediting error in an otherwise correct answer.

TRAP **Treating the tier label as though it were evidence of effect.** Candidates write that an activity is primary prevention and then proceed as though its benefit had been established by the classification. It has not: the tier states when and on whom the intervention acts, and every claim about whether it works is a separate claim requiring its own evidence, its own population and its own outcome. The most common form of this error is asserting that a public awareness activity prevents disorder, when the evidence available concerns knowledge or attitudes and the outcome asserted is incidence.

The two axes answer different questions, when and to whom, and the taxonomy's job in a community psychiatry chapter is to route each application to the place where its evidence is argued in full.

35 · Mental health economics: cost of illness, cost-effectiveness and financing

The argument about what mental health should cost is made whether or not psychiatrists take part in it, and it is settled by people who read budget documents rather than case notes. That is the practical reason this material is examinable: a clinician who cannot say what a service costs, what it saves, who pays and where the money enters the system is unable to argue for it in the only language the argument is conducted in. The topic divides cleanly into a small number of questions that are frequently run together and should not be. What does this illness cost, and to whom. Is this intervention worth its cost compared with the alternative. How does money actually reach a mental-health service. And what happens between an allocation and an expenditure. The classic economic argument for funding a general-hospital liaison service, with its historical figures and the observation that Indian cost-effectiveness evidence for such services is what is missing, is already made in the Consultation-liaison, geriatric and transcultural psychiatry notes and is not repeated or re-declared here. Insurance benefit design, parity and the regulator's mandate belong to the insurance and financial-protection material of these notes.

35.1 Cost of illness, and what the total is and is not

A **cost of illness** study answers a descriptive question: what does this condition consume, across a defined population, over a defined period. It is a burden statement, not an evaluation, and it makes no comparison between courses of action. Its components separate as follows, and the separation matters because different components fall on different payers.

COMPONENT	WHAT IT COUNTS	WHO USUALLY BEARS IT
Direct medical	Consultations, medicines, investigations, admissions, procedures	The health system, the household, or an insurer, in a proportion that varies by country
Direct non-medical	Transport to care, food and lodging for an accompanying relative, informal payments	The household, almost entirely, and it is routinely omitted from health-system costings
Indirect	Lost work and productivity of the patient, and of family members providing care	The household and the wider economy; the largest component for most mental disorders
Intangible	Pain, distress and loss of quality of life not otherwise captured	The patient and family; usually acknowledged and not costed

Two features of a cost of illness total decide whether it can be used at all. The first is the **perspective**: a total computed from the health system's point of view counts what the health system spends, and a total computed from a societal perspective adds the household's costs and the lost production. These give very different numbers for the same disease, and in mental illness the difference is unusually large because the indirect component dominates. A figure quoted without its perspective is uninterpretable.

The second is the counterfactual, and this is where the figure is most often misused. A cost of illness total is not a sum that could be saved. It is what the condition currently consumes, and no intervention removes all of it; some of it would be consumed by whatever the person had instead, and some of it is the cost of the treatment that keeps them well. The only defensible use of the total is to establish scale and to argue that a problem consuming this much deserves attention, and the only way to establish a saving is a comparison between courses of action, which is a different study.

TRAP **Presenting a cost of illness figure as the money a proposed programme would save.** The two answer different questions and the first cannot substitute for the second. A burden total has no comparator, so it cannot support any claim about a difference; and treating the whole burden as recoverable implies an intervention that abolishes the disorder at no cost, which nothing does. If the question is whether a programme is worth funding, the answer requires an economic evaluation and not a burden estimate.

35.2 What India's cost of illness figure is, and why it is not printed here

An examiner may reasonably expect an Indian national figure at this point, and the honest answer is that this chapter does not have one. What was sought was an official India-wide annual cost of illness estimate for mental disorders, with its direct and indirect components separated and a stated price or reference year. The national survey's summary report was consulted, the detailed report's official address returned an error, the ministry's mental-health and programme pages were searched, and indexed-literature and international-agency routes were tried for an Indian economic estimate. No official national estimate meeting that description was obtained.

The absence is worth teaching rather than concealing, and it is not a small one. Without an official national burden estimate, an Indian argument for mental-health investment has to be assembled from international burden estimates, from disability measures whose derivation belongs to the burden material of these notes, and from individual studies with local populations. Each of those is defensible and none of them is what a finance ministry asks for, which is one figure attributable to a national body with a stated year. A candidate who names the construct, says the national estimate is not established, and explains what the argument has to be built from instead has given a better answer than one who quotes a figure whose issuing body they cannot name.

35.3 Economic evaluation, and the four forms it takes

An **economic evaluation** always compares at least two courses of action on both their costs and their consequences. If it compares costs only, it is a costing exercise. If it compares consequences only, it is a clinical study. The forms differ in one respect only, which is how the consequence is valued, and that difference determines what question each can answer.

FORM	HOW THE OUTCOME IS VALUED	THE QUESTION IT ANSWERS	ITS LIMITATION
Cost-minimisation	Not valued; outcomes are assumed equivalent	Which of two equally effective options is cheaper	The equivalence assumption is rarely justified and is often asserted rather than shown
Cost-effectiveness	In natural clinical units, such as symptom-free time or cases detected	What it costs to buy one more unit of this particular outcome	Cannot compare across conditions, because the units differ
Cost-utility	In a combined measure of length and quality of life	What it costs to buy a comparable unit of health, whatever the condition	Depends on how the quality weights were derived and in which population
Cost-benefit	In money, including the health outcome itself	Whether total benefits exceed total costs, in one currency	Requires putting a money value on health, which is contested and rarely done well

The practical significance for mental health is that the middle two dominate, and the choice between them is a choice about who the result must persuade. A **cost-effectiveness** result in natural units speaks to a psychiatrist and to a programme manager who already accept that the condition matters. A cost-utility result is the one that lets mental health compete for a share of a health budget against cardiology and oncology, because it converts the outcome into a currency the whole budget uses. A specialty that only ever produces results in its own units is arguing where nobody else is listening.

35.4 Reading a cost-effectiveness result without being deceived by it

A cost-effectiveness figure carries a set of design choices, each of which can move it substantially, and each of which is a legitimate examination question.

- The perspective, again: a result from the health system's point of view ignores what the household paid to attend, and interventions that shift cost onto families look efficient from inside the system.
- The comparator: an intervention compared with nothing will almost always look good, and the informative comparison is with what the service would otherwise have done.
- The time horizon: benefits that accrue late are missed by a short study, and mental-health interventions with relapse-prevention effects are systematically undervalued by short follow-up.
- Discounting: costs and benefits arriving at different times are not equivalent, and the rate chosen affects any intervention whose costs are immediate and whose benefits are later, which describes most of prevention.
- Uncertainty: a single point estimate hides the range of plausible results, so a defensible evaluation reports how likely the conclusion is to hold rather than asserting it.

An Indian result illustrates the last of these particularly well. In a trial-based economic evaluation conducted in primary health centres in Goa, the programme under evaluation, identified in the trial by the acronym **HAP**, delivered alongside enhanced usual care, had a **58 per cent** chance of being cost-saving from a health system perspective.

Notice the form of that statement, because the form is the teaching. It is not a claim that the programme saved money; it is a statement of how often the intervention came out cheaper across repeated resampling of the paired cost and outcome data. A probability just above the halfway mark is a genuinely useful result and an honest one: it says the intervention is about as likely as not to pay for itself in health-system costs, while delivering its clinical benefit. Reported as a point claim that the programme is cost-saving, the same analysis would be a misrepresentation, and reported as a failure to demonstrate saving it would be a different one.

58 per cent	Not established	Not retrieved	Withheld
CHANCE OF BEING COST-SAVING, HEALTH-SYSTEM PERSPECTIVE, ONE GOA PRIMARY-CARE TRIAL	OFFICIAL INDIA-WIDE ANNUAL COST OF ILLNESS FOR MENTAL DISORDERS	LIKE-FOR-LIKE NATIONAL ALLOCATION AGAINST ACTUAL EXPENDITURE	SCHEME-LEVEL BUDGET LINES NO CLAIM ROW HERE ASSERTS

RULE **A cost or cost-effectiveness figure means nothing without its perspective, its comparator and its time horizon, and a probability of being cost-saving is not a statement that it saved money.** Those four attributes are the whole of the appraisal, and a figure quoted without them is the economic equivalent of a prevalence quoted without its population.

35.5 How money actually reaches a mental-health service

Financing mechanisms are the least glamorous part of this topic and the part that most determines what a district can do. The generic mechanisms are few. General taxation funds a service directly from government revenue. Social or mandatory health insurance funds it from contributions. Voluntary or private insurance funds it from premiums. Out-of-pocket payment funds it from the household at the point of use. External or donor financing funds it from outside

the country. Most systems, India's included, use several at once, and the mix determines who is protected and who is not.

What matters clinically is that these mechanisms differ in when the money is paid relative to when the illness occurs. Pre-paid mechanisms, whether tax or insurance, spread the cost across a population and across time. Out-of-pocket payment concentrates it on the household at the moment it is least able to bear it, which for a chronic mental illness means repeatedly, over years, in a family whose earning capacity the illness has already reduced. That is the whole argument for financial protection, and it is an argument about the timing of payment rather than about its total.

The Indian route by which programme money reaches a state has a specific and consequential shape. Under the **National Health Mission**, money reaches a state against that state's own Programme Implementation Plan. The plan is the state's own proposal of what it intends to do, and the money follows the proposal.

IN INDIA

Because national mission money reaches a state against its own Programme Implementation Plan, the operative lever for anyone trying to fund a mental-health activity in a district is what gets written into that plan before the financial year, not what is argued for afterwards. An activity nobody proposed is an activity nobody funded, and a clinician who wants a post, a training programme or a drug budget has a specific and dated point of entry: the plan-writing cycle. This is the single most actionable piece of financing knowledge in the topic, and it is administrative rather than economic.

35.6 Allocation against expenditure, and the gap between them

A budget line is a permission to spend, not a record of spending, and in mental health the two diverge more than in most programmes. The reasons are structural and are worth knowing because they are what an examiner is testing when they ask why increased allocation has not produced increased service. Money allocated to posts cannot be spent if the posts cannot be recruited to, and specialist mental-health posts are among the hardest to fill. Money allocated late in a financial year cannot be spent before it lapses. Money allocated for capital cannot be moved to recurring costs even when recurring costs are the constraint. And money released against a plan that was written optimistically returns unspent, which then becomes an argument for allocating less next time.

That last mechanism is the important one and it is self-reinforcing. Under-spending is read by a finance authority as evidence of low absorptive capacity, which justifies a smaller allocation, which reduces the capacity further. Breaking that loop requires spending on the things that can actually be delivered, which usually means training, supervision and recurring supply rather than infrastructure, and it requires the plan to be written by somebody who knows what the district can absorb.

PITFALL

Quoting an allocation figure as though it described a service. An allocation is what a document says may be spent; it is not what was released, not what was spent, and not what reached a patient. The chain from allocation to release to expenditure to utilisation loses value at each join, and a claim about service delivery supported only by a budget number has skipped every one of them.

35.7 The financing figures this chapter deliberately does not print

This section is where a chapter is most tempted to supply remembered numbers, and several that would ordinarily appear are declared here as gaps instead. Each is a real construct with a real answer somewhere; none is established by anything this chapter could open and verify, and a figure written from recollection into a financing argument is the kind of error that survives into somebody's answer script.

No like-for-like national total of mental-health allocation and actual expenditure, across the union and all states and union territories with each financial year identified, was obtained. Union budget documents, expenditure documents, the ministry's annual report, the national programme's own page, state plan and record-of-proceedings indexes, and the programme's framework and plan guidelines were all approached. The pieces exist in different documents on different bases and were not reconcilable into one comparable total.

The national programme's own line in the detailed demands for grants could not be reached: the address that would carry it returns an error after a domain migration, a repair attempt returned the same dead address, and no archived snapshot was available. The scheme-level budget lines for the national tele-mental-health programme, across recent actual, budget and revised estimates, and the plan-period allocations to the district programme, are likewise not printed here: those figures exist in budget documents but no claim row in this chapter's ledger asserts them, so they are withheld deliberately rather than reconstructed.

The teaching that survives all of that is the discipline itself, and it is the discipline the whole chapter is built on. An Indian financing figure is a fact only when it carries the document that issued it, the financial year it refers to, whether it is an allocation, a revised estimate or an actual expenditure, and what it is a proportion of. Those four attributes are exactly what makes budget figures so easy to misquote: the same programme in the same year has several different correct numbers attached to it, and quoting the wrong one confidently is indistinguishable, to a reader, from quoting an invented one.

35.8 Financial protection, and what a household actually pays

The last question is what the illness costs the family, and it is the one a clinician meets most directly. Out-of-pocket payment for a chronic mental illness is not a single event but a stream: repeated medicine purchases, repeated travel, repeated days of lost earning for the patient and for whoever accompanies them, and, where the public service is unavailable or unacceptable, private consultation fees. The stream continues during the periods when the person is well, which is exactly when a family is most likely to stop paying it, and non-adherence for cost reasons is often invisible to a clinician who never asks about the transport fare.

The national survey did produce a household expenditure finding for care of a person with a mental disorder, together with an observation about the absence of state or insurance coverage for most families. That finding is a declared gap in this chapter and its figure is not printed, because no claim row here asserts it. The construct is named without its rupee amount, and the reasoning around it stands on its own: an expenditure of this kind, incurred monthly and indefinitely by a household whose earning capacity the illness has reduced, is a financial-protection failure whatever its exact size.

The specific Indian mechanisms of financial protection, meaning the public insurance scheme, the parity requirement for mental illness and the regulator's mandate, are taught in the insurance and financial-protection material of these notes and are not restated here. What belongs to this section is the economic point behind them: financial protection is achieved by moving payment from the point of use to a pre-paid pool, and any scheme is judged on whether the household's payment at the point of use actually falls, not on whether a benefit exists on paper.

> A cost of illness total states scale and can never state a saving; only an economic evaluation with a named perspective, comparator and time horizon can do that, and in India the money reaches a service through the state's own plan.

36 · Digital mental health interventions and apps in the Indian context

The interesting question about a mental-health app in India is almost never whether it works, and almost always whether anybody who needs it can reach it, will keep using it, and has somebody at the other end who acts on what it produces. That is the difference between the product question and the system question, and it is the boundary this section works to. Whether a given digital therapeutic has established efficacy, what trial standards it must meet, how software is classified for regulatory purposes and what it means for such a product to be prescribable are taught in the Recent advances and landmark studies notes. What belongs here is the Indian implementation picture: access and the divide inside it, adoption and adherence, integration with an existing service, equity, governance and programme monitoring. Remote consultation between a patient and a clinician is a service model in its own right and is taught in the telepsychiatry material of these notes; this section covers what the health system does with software rather than what it does with a video link.

36.1 A functional taxonomy, because the technology label explains nothing

Grouping digital tools by their platform, meaning apps against websites against messaging, tells you nothing useful. Grouping them by what function they perform in a service tells you who has to act, what has to exist behind them, and how they fail.

FUNCTION	WHAT IT DOES	WHO MUST ACT FOR IT TO MATTER	HOW IT CHARACTERISTICALLY FAILS
Information and psychoeducation	Explains a condition, a treatment or a service to a public or a patient	Nobody, in the short term; it is a one-way transfer	Written in a register and a language its intended audience does not use
Self-guided intervention	Delivers a structured therapeutic programme with no clinician involved	The user alone, repeatedly, without prompting	Attrition, which is very high without human contact
Guided or blended intervention	Delivers the same programme with a person supporting adherence	A supporter, who need not be a specialist	The support component is dropped first when the budget tightens
Clinician-facing decision and workflow support	Prompts assessment, records outcomes, flags who is not improving	The clinician, and whoever governs the workflow	Added to a workload with no time released, so it is bypassed
Screening and case-finding by a frontline worker	Structures an enquiry the worker is already making	The worker, and the service receiving the positives	Positives are generated with no destination to send them to
Monitoring and follow-up	Collects symptoms, adherence or attendance between contacts	Somebody who reads it and responds	Data accumulate that nobody has been made responsible for

The fourth column is the section in miniature. Every failure listed is organisational rather than technical, and none of them is fixed by a better application. That is why an Indian digital mental-health question is a service question.

RULE **A digital tool is worth what the response behind it is worth.** Software that identifies distress, records a symptom score or flags deterioration changes nothing unless a named person is obliged to act on the output within a defined period. Where no such person exists, the tool has converted a clinical problem into a data problem.

36.2 Access, and the divide inside the population that most needs the service

The **digital divide** is the reason this matters. Digital delivery is routinely presented as an access solution, and it can be, but only against an explicit account of who cannot use it. The entry requirements are specific and each is distributed unevenly in the same direction as the underlying disadvantage.

- A device that the person controls, rather than a household device controlled by somebody else.
- Connectivity that is reliable enough and cheap enough for repeated rather than occasional use.
- The language of the interface, which for a person in distress must be the language they think in and not merely one they can read.

- Enough literacy, and enough familiarity with the conventions of an interface, to navigate it without help.
- Privacy at the moment of use, which is the requirement most often forgotten and the one hardest to meet in a crowded household.

The shared-device problem deserves separate emphasis because it inverts the usual assumption. Where a phone is shared within a household, an application on it is visible to whoever else uses it, and a browsing history, a notification or an installed icon can disclose a person's help-seeking to the very people they may need to keep it from. For a woman in a household where help-seeking is contested, and for an adolescent, a digital route can be less private than a clinic visit rather than more. A service that assumes digital equals confidential has assumed the opposite of what may be true.

> **PITFALL**
>
> Building in one language and adding others later. Language is treated as a translation task to be done after the product works, so the pilot runs in English or in the state's dominant language, the evaluation reports on the population that could use it, and the populations excluded never appear in the denominator. Because the excluded groups are also the underserved ones, the programme's own evaluation will report success while the access gap it was built to close is unchanged.

36.3 The Indian evidence, and what it points to

The strongest Indian signal for digital mental health does not come from a consumer product but from a multifaceted intervention delivered inside a rural primary-care system. In rural primary-health-centre clusters in Haryana and Andhra Pradesh, among adults at high risk of depression or self-harm, those who received the intervention had lower depression scores than those receiving usual care, and nearly every secondary outcome moved in the same direction. The outcomes examined alongside depression included anxiety, remission, access to treatment, knowledge, attitudes, stigma and intended help-seeking.

Two features of that result are worth extracting, because they are what an Indian answer should be built on. The first is that the digital component sat inside a service rather than beside it: the technology supported frontline health workers and primary-care clinicians who were already in contact with the population, which means the entry requirements listed above were largely met by the service and not by the patient. The second is the breadth of the outcome set. An intervention that moves knowledge, attitudes, stigma and intended help-seeking as well as symptoms is acting on the pathway to care and not only on the illness, which is the correct ambition for a population intervention and is not something a standalone self-guided product is designed to do.

IN INDIA

Because the Indian evidence that has actually moved depression scores comes from an intervention delivered through rural primary-care clusters, associated with lower depression scores than usual care, the practical decision for a district is to treat a digital tool as a component bolted into an existing pathway with named workers behind it, not as a service that can be procured and pointed at a population. The corollary is a procurement rule: before adopting any tool, name the worker who will use it, the time released for them to do so, and the clinician who receives what it produces. A tool adopted without those three has been bought rather than implemented.

36.4 Adoption and adherence, which are separate problems

Digital interventions have a characteristic and well-described failure curve, and it has a shape that no other intervention type shares. **Uptake** is easy and cheap, because installing something costs the user almost nothing. **Continued use** collapses quickly, because stopping also costs almost nothing. The result is that the population who complete a self-guided programme is a small and highly selected subset of the population who started it, and it is selected for exactly the characteristics that predict good outcome anyway.

The single most reliable modifier of that curve is human contact. A brief, scheduled contact from a person, who need not be a mental-health specialist and need not deliver any therapy, substantially changes the proportion who continue. That is the entire logic of guided digital intervention, and it also explains why the guided version costs more than the self-guided one and is worth it: the support is not an accessory to the software, it is the component that makes the software reach anybody.

MNEMONIC

Access, Adoption, Adherence, Action

The four gates a digital intervention has to pass, each with its own denominator and its own failure. Who could reach it. Who started it. Who kept going. And who did something with what it produced. A programme reporting only the second gate has reported the easiest number it has.

Who could reach it	Who started	Who continued	Who was acted on
ACCESS, DENOMINATOR IS THE INTENDED POPULATION	ADOPTION, DENOMINATOR IS THOSE WHO COULD REACH IT	ADHERENCE, DENOMINATOR IS THOSE WHO STARTED	ACTION, DENOMINATOR IS THOSE WHO PRODUCED A SIGNAL

36.5 Integration, or why an application is not a service

An application becomes part of a health service at a small number of specific joins, and naming them is the most useful practical content in this topic. Who introduces the tool to the patient, and at which contact. Whether what the tool records enters the clinical record or sits in a separate system nobody opens. Who is responsible for a concerning response, within what period, and by what route. What happens when the person stops using it, since disengagement from a monitoring tool is itself clinical information. And who supports the user when the tool does not work, which in practice is whoever they saw last.

Safety escalation is the join with the highest stakes and the one most often unspecified. A tool that asks about self-harm will receive positive answers, at any hour, from people whose location and identity it may not know. A programme that has not decided in advance what happens then has built a mechanism for collecting disclosures it cannot respond to. The decision is not primarily technical: it is about who is on duty, what they are permitted to do, and what the local escalation route is.

36.6 Governance, data and the questions an Indian programme has to answer for itself

Digital mental-health data are among the most sensitive categories of personal information, and the governance questions are concrete rather than abstract. Where are the data stored and under whose control. Who may access them, and can an employer, an insurer or a family member obtain them. What was the person actually told, and was consent obtained in a language and a form they understood rather than by scrolling past a screen. What happens to the data if the provider is acquired, changes its terms, or closes. And who is accountable when the tool gives advice that turns out to be wrong.

This chapter carries no Indian instrument that settles those questions for mental-health applications, and none is attributed here. That absence should be stated as such rather than filled: a candidate who names the questions, says that a specific national framework for evaluating, endorsing or governing mental-health applications is not something they can attribute, and points to the chapter that carries the regulatory classification of software has answered better than one who names a framework from recollection. The regulatory and product-classification material sits in the Recent advances and landmark studies notes.

One governance point can be made without any instrument, because it follows from the service logic already established. A public programme that recommends a third-party application to a population has taken on responsibility for it, and has done so without any of the assurance machinery it would apply to a medicine or a device. That is a decision worth making explicitly rather than by default, and the minimum discipline is to know who evaluated the tool, against what, and who reviews that judgement when the product changes.

36.7 Monitoring a digital programme, and the measures that flatter it

Digital programmes generate more data about themselves than any other kind of service, and almost all of it measures the wrong thing. Downloads, registrations, page views, sessions and average session duration are all easy to collect, all report on activity, and none has the intended population in its denominator. A programme reporting them is reporting the traffic its own promotion generated.

The measures that answer the question are the four gates named above, each reported against its own denominator, plus the clinical outcome in the population reached. Two further measures are worth insisting on because they are the ones that reveal an equity failure. First, the composition of the reached population compared with the intended one, broken down at least by sex, by rural or urban residence and by language, since a programme can look successful while serving a demographic that was never the target. Second, the proportion of signals generated that received

a response within the intended period, which is the measure that distinguishes a service from a data-collection exercise.

TRAP **Quoting downloads or registrations as evidence of a digital mental-health programme's reach.** A download is an installation event with no denominator, no population and no clinical content. It can be produced by advertising alone, it counts devices rather than people, and it says nothing about whether the application was opened twice. The reach question requires the intended population as a denominator, and the impact question requires an outcome measured in the people who actually used it.

36.8 What this section does not have, and how to answer without it

The figures a resident is most likely to be asked for here are the ones this chapter cannot supply, and the discipline of saying so is worth as much as the figures would be. No official national Indian figure for the reach, uptake or effectiveness of any mental-health application is carried here, and none is invented. No national count of Indian mental-health applications, of their users, or of their clinical outcomes is available to this chapter. No Indian figure for device ownership, connectivity or private access among people with mental illness is held either, which is why the access argument above is made structurally rather than quantified.

Nor is any prevalence figure supplied for online harms among Indian school-going children and adolescents. That figure is frequently quoted, the populations behind the various circulating versions differ, and a figure measured in one country's schoolchildren does not become an Indian figure by being applied to Indian ones. A service designing a response to online harm has to do so without knowing the size of the problem, which is uncomfortable and is the actual position.

There is a further consequence of that position which is worth drawing out, because it is the one a service actually has to act on. Where the size of a problem is unknown, a service cannot size its response to the problem, so it must size the response to what it can sustain and then measure what it meets. That is a different planning logic from the one used for a quantified problem, and it puts the emphasis on building a route that can absorb whatever arrives rather than on estimating demand in advance. Applied to online harm among young people it means an accessible, non-stigmatising point of contact and a workforce trained to receive the presentation, rather than a campaign scaled to a prevalence nobody has established.

What remains, and it is the durable part, is the framework: a digital intervention is characterised by its function rather than its platform, is limited first by access and then by adherence, is worth what the response behind it is worth, and is evaluated on the intended population as a denominator rather than on the traffic it generated. Those propositions hold whichever product is in front of you, and they are what an examiner is testing when they ask about digital mental health in the Indian context.

Access, adoption, adherence and action are separate gates with separate denominators, and the Indian evidence that has moved outcomes came from a tool embedded in a primary-care service rather than from a product handed to a population.

37 · Social determinants of mental health and global mental health movement

Almost every strong predictor of mental ill health in a population is something a psychiatrist cannot prescribe. Income, education, housing, safety at home, work, and the accumulated experience of childhood do more to shape the distribution of disorder in a district than the availability of any treatment does, and they are administered by departments the health service does not control. That is the uncomfortable structure of this topic, and it is what makes it a public-health subject rather than an aetiological one. The examinable content is not a list of hardships; it is the argument that follows from the list, which runs: the causes are largely social, the levers are largely outside health, therefore the health system's role includes evidence, advocacy and intersectoral work, and none of that reduces its obligation to treat the people in front of it. Migration and ethnicity taught as a clinical and epidemiological topic, with its mechanisms, its acculturation framework and its risk estimates, belongs to the Consultation-liaison, geriatric and transcultural psychiatry notes; migration appears here only as one determinant among others, without any of its figures.

37.1 What a social determinant is, and the life-course claim

A **social determinant** is a condition of life that influences health and is itself produced by how a society distributes resources, status and power. Two properties distinguish a determinant from an ordinary risk factor. It is structurally produced, so it is patterned across a population rather than randomly distributed. And it is modifiable by policy rather than by clinical action, which is precisely why it belongs in a chapter about programmes.

The organising proposition for the whole topic is that mental health, and much of the common mental disorder a population carries, is patterned by **the environments people live through**, meaning the social, the economic and the physical together, acting at different stages of a life rather than at one. Every clause of that proposition is doing work. It covers mental health and not only disorder, so the argument is about the distribution of wellbeing and not merely about caseness. It names three kinds of environment, so the determinants are not reducible to income. And the life-course clause is the decisive one: exposures act at different stages, they accumulate, and an exposure in early life can express itself decades later, which means the timing of an intervention matters as much as its content.

MNEMONIC

Cause, Course, Consequence, Contact

The four routes by which a social determinant reaches mental health, and they are separate claims requiring separate evidence. It can raise the chance that disorder occurs. It can worsen the course once it has. It can amplify the social damage the disorder does. And it can obstruct contact with a service, which is the access limb and belongs with the treatment gap material of these notes.

37.2 The determinants, grouped by the lever that would move them

Listing determinants is easy and unilluminating. Grouping them by which sector holds the lever converts the list into a plan, and it is the form an examiner rewards.

DETERMINANT	HOW IT ACTS ON MENTAL HEALTH	WHICH SECTOR HOLDS THE LEVER
Poverty and income insecurity	Chronic stress, debt, inability to sustain treatment, and constrained choice at every decision point	Finance, social protection and employment
Education	Shapes literacy, earning capacity and the ability to navigate services across the whole life course	Education
Adverse experience in childhood	Acts early, accumulates, and expresses itself long after the exposure has ended	Child protection, education, police and social welfare
Violence in the home and gender inequality	Direct exposure, and constrained ability to leave, to earn or to seek help independently	Home affairs, law, and women and child development
Housing and the physical environment	Crowding, insecurity of tenure, noise, heat, sanitation and the absence of private space	Housing, urban development and local government
Work and its conditions	Demand, control, security and reward, and the loss of all of them in unemployment	Labour, and employers
Alcohol and substance availability	Price and access shape consumption, and consumption shapes harm at population scale	Excise, revenue and law enforcement
Discrimination and social exclusion	Chronic exposure to devaluation, and reduced access to every other resource above	Law, education and every service that receives the public

The third column is the point of the table. Not one of these levers sits inside a mental-health budget, and most sit outside health altogether. A district mental-health plan that lists determinants without naming the department that owns each one has produced a description; a plan that names them has produced an addressee.

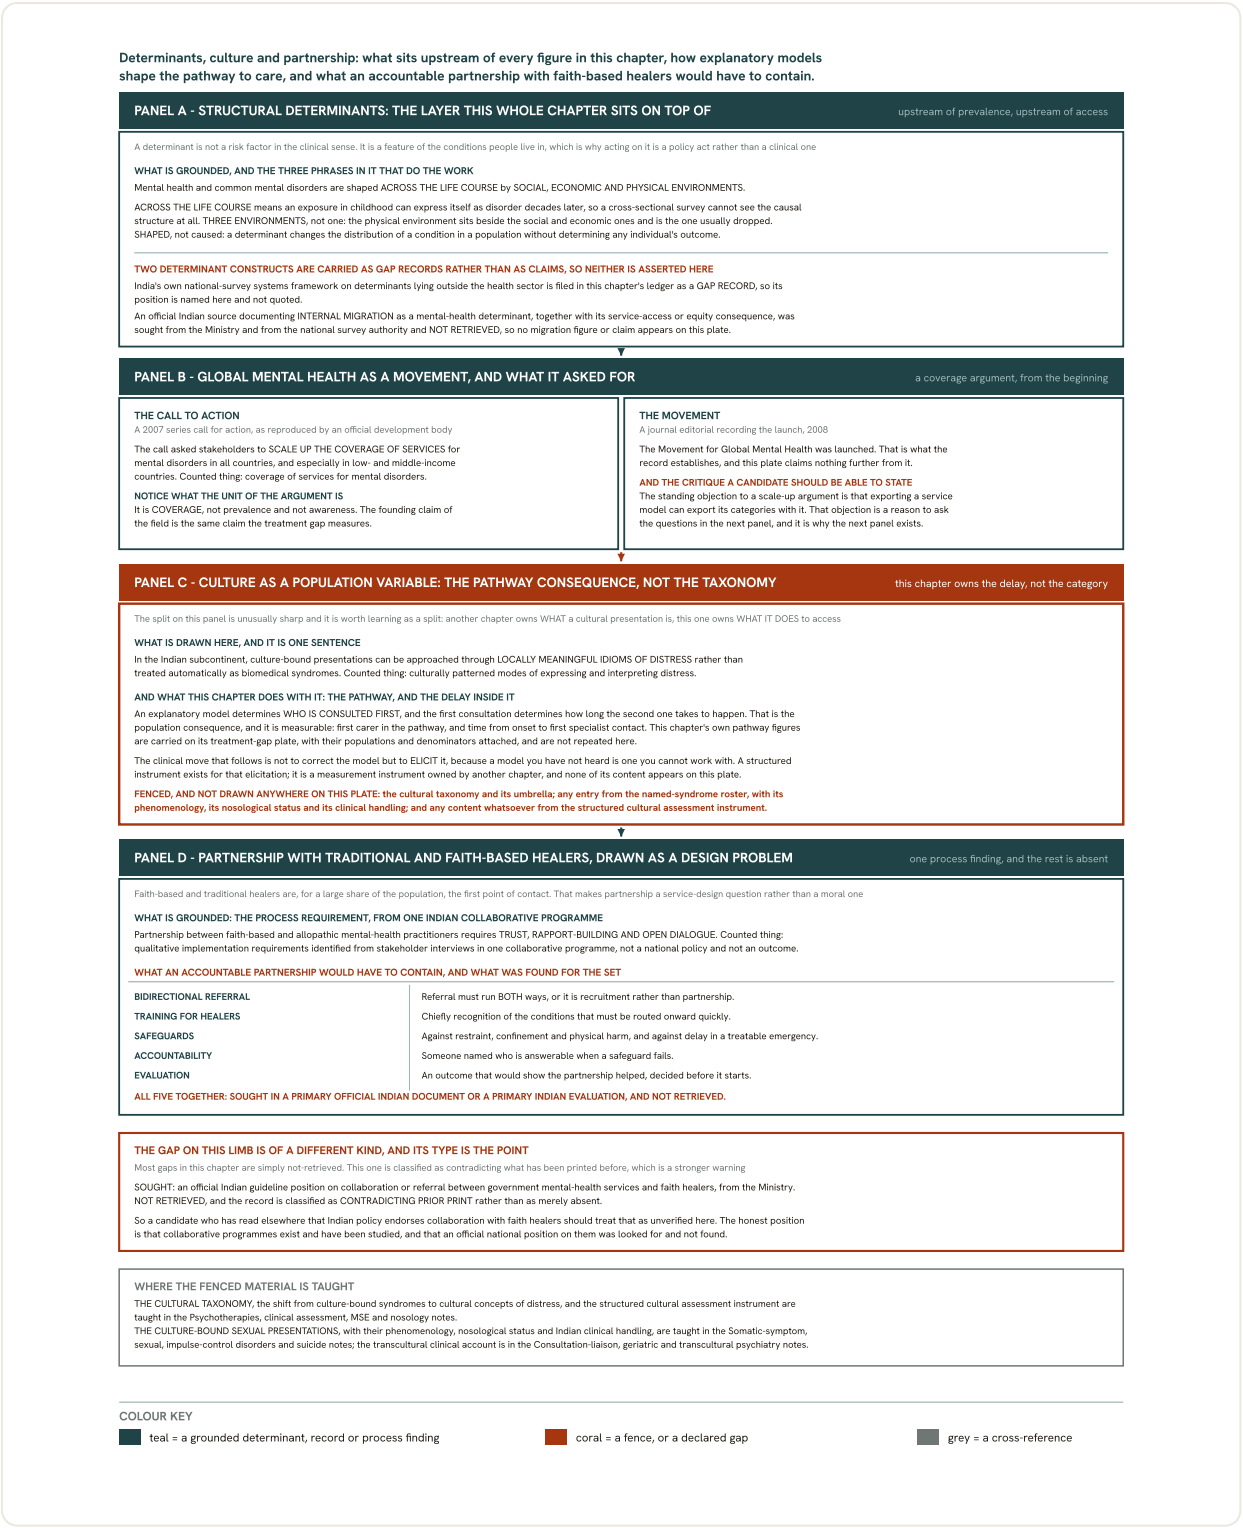


Fig 39.13 Determinants, culture and partnership: the layer upstream of every figure in this chapter, and the two limbs where the evidence runs out. The grounded determinant statement is that mental health and common mental disorders are shaped across the life course by social, economic and physical environments, in which the life-course clause, the three environments and the word shaped are each load-bearing. Global mental health began as a coverage argument: the call to action asked for scale-up of service coverage in all countries, especially low- and middle-income ones, and the Movement for Global Mental Health was launched. Culture is drawn only for its pathway consequence, culture-bound presentations approached through locally meaningful idioms of distress; the taxonomy, the named syndromes and the structured cultural assessment instrument are fenced to their owning chapters. Partnership carries one process finding, that trust, rapport-building and open dialogue are core requirements, while no primary source specifying bidirectional referral, healer training, safeguards, accountability and evaluation together was retrieved and the official Indian position is recorded as contradicting prior print rather than as merely absent.

37.3 The gradient, and what it implies for how a service is targeted

Determinants act as a **gradient** and not as a threshold. It is not the case that disadvantage below some line produces risk and everything above it is equivalent; the relationship runs across the whole range, so a service that concentrates entirely on the worst-off leaves most of the population-attributable burden untouched, because most people are not in the worst-off group.

That produces a design tension every programme in this chapter meets. Targeting the most disadvantaged is efficient per person reached and misses the bulk of cases. Universal provision reaches everyone and dilutes effort where need is lowest. The resolution used in practice is to provide universally while scaling the intensity of provision to need, so that everybody is covered and the most disadvantaged receive more. The principle carries a formal name in the health-inequalities literature, which this section does not attach because no claim row here carries it; the design instruction stands without the label, and a candidate should describe the arrangement rather than name a framework they cannot attribute.

RULE **A determinant is a cause the health system can usually only reach through another sector, and naming the sector is part of naming the determinant.** A statement that poverty causes mental ill health is not yet a public-health proposition. It becomes one when it says which policy instrument, held by which department, would move it, and what the health system's contribution to that argument is.

37.4 The nihilism trap, and the intervention menu that answers it

The commonest misuse of this material is to treat it as a counsel of despair: if the causes are social, clinical work is a distraction, and nothing short of redistribution will help. That inference is wrong in both directions and it is worth being able to dismantle it quickly.

It is wrong about causation, because a determinant raises the probability of disorder in a population without making the disorder in any individual less treatable. It is wrong about action, because the health system has several genuine roles that require no other department's permission. It can produce and publish the evidence that a determinant is acting, which is often the only quantitative case anybody has. It can advocate on the basis of that evidence in the forums where policy is set. It can design its own services so that they do not compound disadvantage, meaning free at the point of use, near where people live, in the language they speak, at hours that do not cost a day's wages. It can deliver interventions that buffer a determinant it cannot remove, such as support to families under strain. And it can treat, properly and without apology, the people the determinants have already made ill.

Set out as a sequence, those roles answer the objection completely, and they are worth being able to recite because they convert a general complaint into a work plan.

- Measure the determinant in the local population, since the health system frequently holds the only routine data on it.
- Publish and advocate, in the forum where the relevant policy is actually decided rather than in a clinical one.

- Remove the barriers the service itself creates, meaning cost, distance, language and appointment hours.
- Deliver the buffering interventions the health system can deliver alone, such as support to families under strain.
- Treat, at the standard the evidence supports, without letting social causation become a reason to expect less.

37.5 Intersectoral action, and what this chapter cannot attribute

Intersectoral action is the practical form of everything above, and it fails in predictable ways. It fails when it is convened without a shared measure, so each department reports its own activity and nobody can say whether anything changed. It fails when the health sector arrives with a request rather than with evidence and an offer. It fails when the benefits accrue to one department and the costs to another, which is the ordinary case. And it fails when it exists only at the level of a committee, with no counterpart at the district level where services actually meet.

An Indian national statement to this effect would be the natural anchor for this section, and it is declared here as a gap rather than supplied. The construct sought was the national survey's own systems framework proposition, that the determinants of mental health lie outside the health sector and that systems for care are incomplete if health-related sectors are not included. That proposition is not established here from a source this chapter opened and verified, so it is named as a construct and not asserted as a finding or attributed to a body. The reasoning above stands on its own as service reasoning; the national attribution does not.

IN INDIA

Because mental health and common mental disorders are shaped by social, economic and physical environments acting at different stages of life, a district mental-health plan that contains only clinical activities has, by construction, left its largest levers unaddressed. The practical decision is to require that the plan name, for each determinant it identifies, the department that holds the lever and the officer who will be approached, so that intersectoral intent becomes a set of addressees rather than a paragraph. A determinant with no named counterpart in another department will not be acted on, and writing it down is what converts the health system's evidence into somebody else's agenda item.

37.6 Migration, taken as one determinant among others

Internal migration belongs in this list because moving, whether across a border or within a country, rearranges almost every other determinant at once: income, housing, language, social network, entitlement and legal standing all change together, usually abruptly. That is why it behaves as a determinant rather than as a single exposure.

The boundary here is strict and is stated so a reader knows where to go. Migration and ethnicity taught as a clinical and epidemiological topic, with the competing mechanisms proposed to explain differences in rates, the framework used to describe adaptation after migration, and the estimates attached to any of it, are taught in the Consultation-liaison, geriatric and transcultural

psychiatry notes. This section prints no risk estimate for migrants in any generation, no Indian migration rate, and no estimate of any protective effect, and it attributes no position on migration to any Indian guideline.

What remains, and it is genuinely this section's, is the Indian internal picture as a service and equity problem. Internal migration within India moves large numbers of people between states and from rural to urban areas, and the service consequence is structural: entitlements, registrations and continuing prescriptions are commonly attached to a place of origin, while the person is now somewhere else. A worker who moves for seasonal or construction work loses continuity of any long-term treatment as a matter of administration rather than of clinical decision, and returns to their district when unwell. Any service designed for a settled population will systematically under-serve them.

The Indian evidence for that argument is declared as a gap. What was sought was an official Indian source documenting internal migration as a mental-health determinant together with its service-access or equity consequence, without reproducing the clinical and epidemiological migration material owned elsewhere. The national survey routes, ministry pages and international agency material approached returned general determinants or mentions of migration among proposed determinants, but no Indian estimate or official statement of the service consequence was retrieved. The structural argument above is therefore offered as reasoning about how entitlement systems work, and no Indian figure is attached to it.

PITFALL

Attaching a risk multiple to migration. The figures that circulate for psychotic illness in migrant populations come from particular host countries, particular migrant groups and particular ascertainment methods, and they are the subject of an active methodological argument taught in the chapter that owns the topic. Importing one into an Indian internal-migration argument compounds two errors at once: it transfers an estimate across populations that differ in almost every relevant respect, and it converts a service and equity argument into an aetiological claim that this chapter has no standing to make.

37.7 The global mental health movement, and what it actually argued

The **global mental health movement** is best understood as an argument rather than an organisation, and its central proposition is straightforward. A large share of the world's burden of mental disorder falls in countries with the fewest specialists, most people with a disorder in those countries receive no treatment, effective interventions exist and can be delivered by non-specialists, therefore the ethical and practical priority is to scale up coverage rather than to wait for specialist workforces to grow.

The movement's defining call, issued in **2007**, asked its readers to increase the coverage of services for mental disorders in every country and above all in low-income and middle-income ones. The wording contains the argument in miniature: what is demanded is *coverage*, which is a population measure with a denominator, and the demand is addressed to every country with a

particular emphasis rather than to poor countries alone. Shortly afterwards the Movement for Global Mental Health was launched as a network to carry that agenda.

2007	Coverage	Non-specialists	Not established here
CALL TO SCALE UP SERVICE COVERAGE, ALL COUNTRIES, ESPECIALLY LOW- AND MIDDLE-INCOME	THE MEASURE THE CALL DEMANDED, NOT SPECIALIST NUMBERS	THE WORKFORCE THE ARGUMENT RESTS ON	ANY INDIAN COVERAGE FIGURE FOR THIS ARGUMENT

The movement has standing objections, and a candidate who can state them fairly is demonstrating more than one who cannot. The first is that exporting a diagnostic and treatment framework developed in one part of the world risks displacing local understandings of distress and local ways of responding to it. The second is that scaling up treatment can absorb the entire agenda and leave the determinants untouched, which is the tension this whole section is about. The third is methodological: an intervention shown to work in one setting is not thereby shown to work in another, and the transportability of evidence is an empirical question rather than an assumption. The fourth is about who sets the agenda, since priorities, funding and publication have historically been concentrated in the countries with the fewest patients.

37.8 What the argument leaves an Indian service with

India occupies an unusual position in this debate, because it is simultaneously one of the settings the movement was arguing about and one of the places that produced the evidence the movement rests on. That gives an Indian clinician a stronger footing than the debate's usual framing allows: the question is not whether an external agenda should be accepted, but which parts of a largely home-grown evidence base a service should implement, and with what adaptation.

Three implications are worth carrying away. Coverage, not capacity, is the measure the argument is conducted in, so a service arguing for resources should be able to state what proportion of its intended population it reaches. Non-specialist delivery with specialist supervision is the mechanism the argument depends on, which is why it recurs in almost every section of this chapter. And the objections above are not reasons to abandon the programme; they are the design constraints it has to satisfy, meaning that a scaled-up service must be culturally intelligible, must be evaluated locally rather than assumed to transfer, and must not be allowed to become the whole of a country's mental-health policy while the determinants that fill it go unaddressed.

TRAP **Treating the social determinants as an alternative to services, or services as an alternative to the determinants.** Candidates pick one and argue against the other, and both arguments are weak. Determinants shape how many people become unwell; services determine what happens to those who do. Neither substitutes for the other, and the examinable position is that a health system owes evidence and advocacy on the first while delivering coverage on the second.

The determinants of mental health are largely produced outside the health sector and reached through other departments, and the global argument that follows is about coverage of a population rather than about the number of specialists a country has.

Culture, global mental health and traditional healing

38 · Cultural psychiatry, idioms of distress, culture-bound syndromes and explanatory models in the Indian context

Culture belongs in a public-health chapter because it decides where a population goes first, how long it stays there, and what it is holding when it finally arrives at a health service. That is a service fact with denominators and waiting times attached, and it is distinct from the cultural encounter at the bedside. The definitions of the cultural constructs, the umbrella that groups them and the classic roster of named syndromes are taught in the Psychotherapies, clinical assessment, MSE and nosology notes; the clinical encounter with a culture-named presentation, the negotiation between two accounts of one complaint and the placement of a cultural formulation inside a working assessment are taught in the Consultation-liaison, geriatric and transcultural psychiatry notes. This section takes none of that. It takes the population consequence: which door a community uses first, what the detour costs, what a general health service has to be able to receive, and what a service must build for any of this to change.

38.1 The pathway, read as a service problem rather than a curiosity

A **pathway to care** is the sequence of people a person consults between first noticing something is wrong and arriving at the service that will eventually treat them. Read at population scale it has three properties that matter to a planner and are invisible in the consulting room.

The first is that psychiatric services are almost never the first stop, and in many communities they are not the second or third either. The first consultation is usually inside the family, then with whoever the community treats as the appropriate authority for that kind of trouble, and then with a general practitioner or a general hospital. That is not a failure of the population; it is a rational sequence given how the problem was understood at the outset.

The second is that each step consumes time and money, and both are spent before any effective treatment starts. The cost is therefore paid twice over: once in the direct expenditure on consultations that did not help, and once in the untreated duration of the illness, which for several conditions is itself associated with worse outcome.

The third, and the one that turns this from description into planning, is that the pathway is modifiable. Where the population goes first is determined by what the community understands the problem to be, and by which doors are physically and socially available. Both can be worked on. A service that treats the pathway as fixed has accepted a delay it could have shortened.

RULE **The pathway to care is a property of the service system as much as of the culture, and it is measurable.** Where a population goes first is determined jointly by what it believes about the problem and by which doors are open, affordable and near. A service can change the second of those directly and the first only slowly, which is why pathway work begins with placing the service where the population already goes.

38.2 The detour, and what it actually costs a service

The **interval between onset and effective treatment** is the quantity that summarises all of this, and it is the number a service should want most and usually does not have. What can be said without any figure is what the delay is composed of, because each component has a different remedy.

STAGE OF THE PATHWAY	WHAT HOLDS THE PERSON THERE	THE LEVER A SERVICE ACTUALLY HAS
Inside the family, before any consultation	The problem is understood as something other than illness, or as something the family should manage	Population-facing explanation in the community's own vocabulary, which is the literacy and awareness work of these notes
With a community or faith-based helper	That helper is the recognised authority for this kind of trouble, is nearby, and is trusted	Partnership and referral arrangements, developed in the traditional and faith-based care material of these notes
With a general practitioner or general hospital	The complaint presented is one the general service can respond to on its own terms	Recognition capacity in general services, which is the primary-care integration work of these notes
Referred but not arrived	Distance, cost, stigma, or a referral with no named destination	Where the service is placed, what it costs, and whether the referral names a person and a date

Notice that only the first row is about belief. The remaining three are about proximity, trust, recognition and the mechanics of referral, all of which are ordinary service-design variables. It is a persistent error to treat the whole of the pathway delay as a cultural problem requiring cultural work, when most of its length is produced by where the service is and how hard it is to enter.

38.3 Idioms and explanatory models as population phenomena

A population does not present its distress in the categories a service uses to classify it. It presents in the vocabulary it already has, which may be bodily, social, moral or spiritual, and which is coherent within the community's own understanding of what has gone wrong. The taxonomy that groups these forms of expression, its definitions and the roster of named presentations that belong to it are taught in the Psychotherapies, clinical assessment, MSE and nosology notes, and the umbrella is named here once as the frame this argument rests on, without reproducing any part of it.

The stance this chapter takes toward those presentations is a service stance, and it has a source. In the Indian subcontinent, culture-bound presentations are more productively received as **locally meaningful idioms of distress** than as biomedical syndromes to be translated automatically at the door. Read as a service instruction rather than as a nosological position, that has concrete consequences.

- The presentation that arrives is data about the population's understanding, not noise to be filtered before the real assessment starts.
- A service that only recognises complaints already framed in its own categories will systematically under-detect, because the population is not obliged to speak in those categories.
- Population-facing communication that ignores the local vocabulary is speaking a language the audience does not use to think about the problem.
- An idiom carries information about severity and about what the person expects to happen next, which is why receiving it accurately affects what can be offered.

Two Indian presentations, dhat and koro, are the standard examples of a complaint whose form is shaped by a shared explanatory model, and they are named here only as that, examples of how a population's understanding routes a person to one door rather than another. Their phenomenology, their nosological status and their clinical handling are taught in the Somatic-symptom, sexual, impulse-control disorders and suicide notes and are not touched here.

IN INDIA

Because culture-bound presentations in this subcontinent are better received as locally meaningful idioms of distress than converted into a biomedical label at the door, an Indian outpatient service should build the recording of the person's own account of the problem into its first contact rather than leave it to whoever has time. The account is what the eventual explanation, the treatment plan and any psychoeducation have to connect to; a service that records only its own formulation has kept the half of the encounter it already understood and discarded the half that determines whether what it offers makes sense to the person receiving it.

38.4 Recognition in general health services, which is where most of it lands

The great majority of culturally framed distress presents to general health services and not to psychiatric ones, which makes **recognition capacity** in those services the single largest determinant of whether it is ever addressed. The design problem is specific: a general practitioner working under time pressure will act on the complaint as presented, and a complaint presented in bodily terms will generate an investigation rather than an enquiry.

The levers here are unremarkable and effective. Training aimed at recognising the common presenting forms in that population rather than at diagnostic subtlety. A short structured enquiry the clinician can actually complete in the time available. A referral destination that is named, near and reachable, since a referral into an abstraction is not a referral. And feedback to the referring clinician about what happened, which is the only mechanism that improves recognition over time and is almost universally absent.

PITFALL

Designing the service around the assumption that the health system is the first door. Outreach, awareness material and referral arrangements are all built as though the population's starting point were a health facility, so they address people who have already found the system. The population that has not is reached by none of them, and it is exactly the population whose pathway is longest. Finding out which door a community actually uses first, in that district, is a prior question and not a research luxury.

38.5 What a service must build for a cultural formulation to be usable at all

It is one thing to accept that a person's own account of their problem should inform their care, and another to run a service in which that happens reliably. The implementation questions are three, and they are the ones a programme has to answer.

The first is who is trained. If eliciting an account is treated as a specialist skill, it will happen only where specialists are, which is the smallest part of the system and the last stop on the pathway. If it is treated as a competence for whoever holds the first clinical contact, it happens where the population is, which requires a short, teachable version and a supervision structure behind it.

The second is at which contact it is applied. Applying it at every contact is unaffordable at scale; applying it at none makes the commitment decorative. The workable answer is to fix it to a defined point, usually the first assessment and any point at which the plan is not working, and to record the outcome where the next clinician will see it.

The third is what it costs in consultation time, and this is the question that decides whether the policy survives contact with a busy outpatient department. An enquiry that adds materially to every consultation in a clinic seeing very large numbers will be abandoned within weeks and will then be reported as having been implemented. The honest planning position is that time is the binding constraint, that the enquiry must be scoped to the time actually available, and that if the time does not exist the service should say so rather than mandate something it cannot deliver.

What such an enquiry is for is modest and specific. **An explanatory-model assessment** is aimed at three things a service otherwise never learns: what the problem means to the person, where they think help might come from, and what they expect a service to do. Each of those is directly actionable, which is why the enquiry is a service investment rather than an ethnographic one: the first tells you what the explanation must connect to, the second tells you who else is already involved, and the third tells you what will be experienced as help. The instrument built on this approach, its structure and its content are taught in the Psychotherapies, clinical assessment, MSE and nosology notes, and where the enquiry belongs inside a working clinical assessment is taught in the Consultation-liaison, geriatric and transcultural psychiatry notes.

MNEMONIC

Place, Partner, Practitioner, Proof

The four decisions that convert a cultural argument into a service change. Place the service where the population's pathway already runs. Partner with whoever holds the door it uses first. Train the practitioner who takes the first clinical contact, not only the specialist at the end. And prove it, by measuring where people came from and how long they took to arrive.

38.6 Measuring a pathway, which almost no service does

The whole of this argument becomes tractable the moment a service records where its patients came from. That is a small addition to a first-assessment record and it produces the only data that can tell a district where its access problem actually is.

First contact	Number of steps	Interval to treatment	Referral source
WHO THE PERSON CONSULTED BEFORE ANY HEALTH SERVICE	HOW MANY HELPERS WERE SEEN BEFORE ARRIVAL	TIME FROM ONSET TO EFFECTIVE TREATMENT, PER PATIENT	WHICH DOOR ACTUALLY DELIVERS PATIENTS TO THIS SERVICE

The last of the four repays a moment's thought, because a service usually believes it already knows the answer and is usually wrong. A department that runs an outreach programme, a helpline and a referral arrangement with a general hospital will assume its patients arrive in some rough proportion across the three. Counted, the distribution is almost always lopsided, and the lopsidedness is the finding: one route is carrying the service and the others are consuming staff time to deliver very few people. That is an actionable result obtainable from a single field on a form.

Those four quantities, collected on a service's own attenders, answer questions no national dataset can. If the first contact is overwhelmingly one kind of helper, the partnership priority is settled. If the number of steps is large but the interval is short, the problem is misdirection rather than delay. If the interval is long, the question becomes which step it accumulated at. And if one referral source supplies most arrivals, the service knows both where to invest and where it is currently invisible.

The important limitation of these measures is that they describe the people who arrived, which makes them a **clinic-attender sample** and not a population one. They say nothing about those who did not, who are the population of interest and are systematically different. Pathway data from a clinic are therefore a starting point for a district question and not an answer to it, and a service that forgets this will optimise for the population it already has.

38.7 What is not established here, and how to answer without it

This section prints no prevalence for any culture-named presentation, and that is a deliberate position rather than an omission. The figures that circulate for such presentations come from clinic samples of very different kinds, use different criteria, and describe populations that are not comparable; a proportion measured among attenders at one specialist service is not a population

rate and cannot be reported as one. The chapter that owns the clinical material makes the same refusal in print, and this chapter does not supply what that one declined to.

Nor does this section carry an Indian national figure for pathway delay, for the proportion of people whose first contact is a non-medical helper, or for the length of the detour. No such national quantity is held here. The consequence for an answer is straightforward and worth practising: describe the structure of the pathway, name the stages and the levers, say that the national quantities are not established, and point out that a service can generate the local versions itself with four fields on its own assessment form. That is a stronger answer than a remembered percentage, and it is also the more useful one, because a district acts on its own numbers and not on a national average.

TRAP **Answering a question about culture and access with a list of named syndromes.** The roster is memorable, it belongs to another part of the syllabus, and reciting it answers a question about nosology rather than the one that was asked. The question here is about where a population goes, why, what that costs, and what a service does about it. A candidate who names the umbrella once, sends the examiner to the chapter that defines it, and then talks about first contact, detour, recognition and referral has answered the question in front of them.

Culture sets which door opens first; the service decides how far away the second one is, and most of a pathway's length is built from distance, cost and recognition rather than from belief.

39 · Traditional healers, faith-based care and integration with mental health services in India

For a large part of the Indian population the first person consulted about a mental-health problem is not a clinician, and no amount of service planning changes that by wishing. The public-health question that follows is narrow and practical: given that a first contact exists which the health system does not control, what arrangement should the health system have with it. That is a design question about partnership, and it is what this section teaches. The clinical encounter with a patient who has been to a healer, the single measured Indian evaluation of a healing site and what it can and cannot bear, the access-against-acceptability distinction and the clinical posture of collaboration without endorsement are taught in the Consultation-liaison, geriatric and transcultural psychiatry notes. Traditional and indigenous resources considered as an Indian service model in their own right are developed elsewhere in these notes. What belongs here is the architecture of a partnership: what it is for, how referral runs in both directions, what training can honestly attempt, what safeguards are non-negotiable, and how the whole arrangement is held to account.

39.1 The premise, stated without sentimentality

Two propositions have to be held together, and most weak answers drop one of them. The first is that faith-based and traditional healing is not an evidence-based treatment for mental disorder; that is the standing position and this chapter makes no stronger claim for it. The second is that a

great many people consult such practitioners first, that this is a coherent choice given how the problem is understood locally, and that they are frequently nearer, cheaper, more available and more socially acceptable than any psychiatric service.

Held together, those two propositions generate the entire design logic. The case for a partnership is not therapeutic. It rests on **access**, meaning that the traditional healer is a point of contact with a population the service does not otherwise reach, and on **harm reduction**, meaning that whatever happens at that contact can be made safer and shorter. A service that argues for partnership on the grounds that healing works has made a claim it cannot support. A service that refuses partnership on the grounds that healing does not work has declined a route to a population that is going there regardless.

RULE **Partnership with traditional and faith-based practitioners is justified by access and by harm reduction, not by therapeutic efficacy.** Keeping that distinction visible is what allows a clinician to collaborate without endorsing, and it is also what stops the arrangement drifting into an implicit claim that the health system does not intend and cannot defend.

39.2 What a partnership is actually for

Stated as design intentions rather than as demonstrated effects, a partnership is built to achieve a small number of specific things, and naming them is what stops the arrangement becoming a gesture.

- Earlier arrival: a person who is going to reach the health service eventually reaches it sooner, because the healer routes them rather than retaining them.
- Reduced harm at the first contact, meaning that whatever practices carry risk are not applied to people with mental illness.
- Continuity: treatment already started is not interrupted while the person is under a healer's care, which is the commonest concrete damage the health system actually sees.
- Reach in the other direction: the healer becomes a route by which information, and sometimes services, reach a population that does not attend clinics.
- Reduced conflict for the patient, who is otherwise required to choose between two authorities and frequently resolves it by concealing one from the other.

The last of those is underrated and is the one clinicians can verify from their own practice. Where the two systems are understood to be in opposition, patients conceal their use of one, which means the psychiatrist does not know what else is being taken or done, and the healer does not know a medicine has been stopped. A working partnership does not require either side to agree with the other; it requires only that the patient can be honest with both.

39.3 Bidirectional referral, and why the second direction is the hard one

Bidirectional referral is the structural core of the arrangement, and its two halves are not equally difficult. Referral in the obvious direction is easy to describe and hard to obtain: the healer identifies a person who needs medical care and sends them. What that requires is not goodwill but three specific things. A recognisable trigger, meaning a short and unambiguous set

of circumstances in which the healer will refer, rather than a request to exercise clinical judgement. A named destination that is reachable, since a referral into an institution is not a referral. And feedback, because a healer who refers somebody and never learns what happened has been given no reason to refer the next one.

The second direction is the one that decides whether the arrangement survives, and it is the one health services find hardest to accept. A partnership in which the health service receives and the healer only sends is not a partnership; it is a recruitment channel, and it will be recognised as such. What the health service can legitimately send in the other direction is not patients for healing but the things the healer's role actually needs: information about mental illness in usable form, a person to ask when uncertain, acknowledgement of the role the healer plays in the community, and in some arrangements a physical presence at the healing site so that people who will not travel to a clinic can be seen.

PITFALL

Designing the referral in one direction and calling it collaboration. The health service writes a protocol that tells healers when to send people to the clinic, delivers a training session about it, and reports a partnership. Nothing flows the other way, no relationship is built, and the arrangement has no value to the person being asked to change their practice. It lasts as long as the project that funded it, and it teaches the healer community that the health system's interest is extractive.

39.4 Training, and what it can honestly attempt

Healer training is the component most often proposed and most often over-specified. It cannot convert a healer into a health worker, and attempting it fails in both directions: the healer's own practice is what gives them standing in the community, and a curriculum that treats it as an obstacle to be corrected will end the relationship.

What training can attempt is narrow and useful. Recognition of a small number of presentations that need medical attention urgently, described in terms the healer already uses rather than in diagnostic categories. Awareness that certain conditions respond to treatment, which is frequently new information and is the single most persuasive content. The specific harms to avoid, stated as practices rather than as principles. And the mechanics of referral, meaning who to contact, how, and what to say.

Two design points follow from experience with any non-specialist cadre and apply with more force here. Training without continuing contact decays, so a single workshop is an event rather than a programme. And training delivered in a register that positions the trainer above the trainee will not be received, whatever its content, which is why who delivers it matters as much as what it contains.

39.5 Safeguards, which are the floor and not a negotiating position

Everything above is conditional on a **safeguarding floor**, and this is the part of the topic where a clinician's obligation is unambiguous. Practices that restrain, confine, deprive of food or water, inflict physical harm, or hold a person against their will are unacceptable irrespective of the

setting in which they occur or the intention behind them, and a partnership that tolerates them is not a partnership but a complicity. That people with mental illness have been chained and confined in domestic, custodial and faith-based settings in India is documented; how widespread or prolonged the practice has been is a separate question that this chapter attaches no figure to, and none should be quoted without its source, its reference year and its denominator. The statutory prohibition on chaining is legislation and is taught with the rest of the mental health law.

The practical safeguarding content of a partnership is therefore a short list of conditions that are stated at the outset rather than discovered later.

ELEMENT OF THE PARTNERSHIP	WHAT IT REQUIRES TO BE REAL	WHAT IT MUST NEVER BECOME
Bidirectional referral	A recognisable trigger, a named reachable destination, and feedback in both directions	A one-way extraction channel with a protocol attached
Training	Recognition, the treatability message, the harms to avoid, and referral mechanics, with continuing contact	An attempt to convert the healer into a health worker, or a single workshop counted as a programme
Safeguards	An explicit floor on restraint, confinement and deprivation, agreed before the arrangement starts	A matter for later negotiation, or something the health service declines to raise for fear of losing access
Accountability	A named person on each side, and an agreed route for raising a concern	A memorandum with no individual answerable for anything in it
Evaluation	Agreed measures collected from the start, including referral counts and what happened to those referred	A report of activities delivered, with no denominator and no outcome

TRAP **Attributing a partnership policy, a referral protocol or a training scheme for faith healers to an Indian national guideline.** This is the specific error to avoid on this topic, because the sentence writes itself and sounds authoritative. No official Indian guideline position on collaboration or referral between government mental-health services and faith healers is established by anything this chapter has opened. Everything in the design above is a proposal with its reasoning attached, and it should be presented as such. A candidate who says that the design is a reasoned proposal rather than a national policy has answered accurately; one who attributes it to a ministry has invented an instrument.

39.6 Accountability, and the difference between a partner and a supplier

A partnership that nobody is answerable for is a document, and documents do not refer patients. Accountability here means something narrower than it does in a formal service contract, because one side is not employed by anybody and cannot be performance-managed. What it can mean, and what is worth insisting on, is that a named person on each side is the point

of contact, that there is an agreed way of raising a concern in either direction, and that concerns raised are actually answered.

The asymmetry of power in the arrangement deserves stating, because it shapes what accountability can realistically look like. The health service arrives with institutional standing, funding and the implicit authority of the state; the healer arrives with community standing and nothing else that a bureaucracy recognises. In that situation an accountability mechanism designed by the health service alone will read as supervision, and supervision is not what was agreed. The workable version is reciprocal: both sides can raise a concern, both sides answer, and the arrangement can be ended by either.

The measures that matter follow from the purposes named earlier and are few. How many people were referred in each direction over a stated period. What proportion of those referred to the health service actually arrived, which is the number that tests whether the destination was real. What happened to them, in outcome terms the healer can understand and will care about. And whether any safeguarding concern arose, how it was raised and what followed. A programme reporting only workshops held and materials distributed has reported its own activity and nothing about the partnership.

39.7 Making it work in practice

The single grounded statement this chapter can make about how such collaboration actually proceeds comes from an Indian programme in Gujarat that brought faith-based and allopathic practitioners together, and it is about process rather than outcome. Collaboration of this kind **requires trust, rapport-building and open dialogue**, and it is challenging.

That is a modest-sounding finding and it is the most operationally useful thing in the section, because it identifies where the cost of the programme actually falls. The expensive component is not the training curriculum, the protocol or the printed material; it is the time somebody spends building a relationship with people who have no obligation to cooperate and some reason to be wary. That time is invisible in a project plan, is the first thing cut when the timeline tightens, and is what the whole arrangement rests on.

IN INDIA

Because collaboration between faith-based and allopathic practitioners requires trust, rapport-building and open dialogue, an Indian district beginning this work should budget the relationship-building period explicitly and assign it to a named person before any protocol is written. Opening with a memorandum, a training calendar or a referral form inverts the sequence: those instruments are the record of an agreement that already exists, and issuing them first is how a partnership becomes a document nobody honours.

MNEMONIC

Two doors, one floor, one ledger

The shape of a defensible partnership. Referral runs through two doors, not one. A safeguarding floor is agreed before anything else and is not negotiable afterwards. And one ledger records what actually happened, meaning who was referred, in which direction, and what became of them.

39.8 What is not established, and how to answer without it

This is a topic on which confident-sounding claims are unusually easy to produce, so the position is stated plainly.

Not established	Not retrieved	Not quantified here	Process, not outcome
OFFICIAL INDIAN GUIDELINE POSITION ON COLLABORATION WITH FAITH HEALERS	INDIAN DOCUMENT SPECIFYING ALL FIVE PARTNERSHIP ELEMENTS	HOW OFTEN OR HOW LONG THE RELIGIOUS PATHWAY IS TRAVELLED	WHAT THE ONE GROUNDED INDIAN STATEMENT DESCRIBES

No official Indian guideline position on collaboration or referral between government mental-health services and faith healers is available to this chapter. What was sought was exactly such a position, in national operational guidance for mental, neurological and substance-use care; it was not established, and the ownership rule for this material is explicit that partnership design and referral must not be attributed to an Indian guideline. The design set out above is therefore a reasoned proposal, and its authority is the reasoning, not a ministry.

No primary official Indian document or primary Indian evaluation specifying all five partnership elements together, meaning bidirectional referral, healer training, safeguards, accountability and evaluation, was retrieved either. Several routes were attempted. The consequence is that the five-element structure above is offered as a design framework assembled from the requirements of the task, not as a transcription of an existing Indian model, and it should be presented that way.

No national quantification of the religious or traditional pathway is supplied here: how often it is used, and for how long before psychiatric contact, is not something this chapter attaches a number to. The single measured Indian evaluation of a healing site, its denominators, its design limitations and the reasons its result cannot be generalised are taught in the Consultation-liaison, geriatric and transcultural psychiatry notes, and its figure is not reproduced here.

What survives all of that is the durable content, and it is what an examiner is testing. A partnership is justified by access and harm reduction rather than by efficacy. It runs in both directions or it does not run. It rests on a safeguarding floor stated before anything else. It is held by named people rather than by a document. And it is evaluated on referrals and what happened to those referred, not on workshops delivered. Those propositions are defensible without a single figure, which is fortunate, because the figures do not exist.

Collaborate for access and for harm reduction, never on a claim of efficacy; agree the safeguarding floor first, make the referral run both ways, and name the people who are answerable for it.

Consolidate and apply

40 · Worked vignettes

A vignette in this territory is never really about the patient in it; it is about the part of the system that patient is standing in. Stems that carry marks here rarely ask for a diagnosis or a drug. They put an ordinary clinical picture in front of the candidate and then plant one further fact: a distance, a vacant post, a medicine that ran out between visits, a family that consulted somebody else first, a figure quoted at a district meeting. That further fact is the **service variable**, and the value of the stem lies entirely in whether the candidate reasons from it or reads past it into a management plan. This closing section introduces no new programme and no new figure. It rehearses the reasoning assembled across this chapter as a reading skill, working each stem from the service variable outwards to the population it belongs to.

The discipline that earns marks is refusing to answer a service question with a clinical one. A woman who cannot collect her medicine because the facility that stocks it is a bus journey away is not asking which agent suits her; she is asking which tier of the system should be holding her follow-up. Where a stem turns on the statute itself, on the wording of a right, on capacity or consent, or on the procedure that follows a death in custody, it has left this territory: those belong to the Mental health legislation and forensic psychiatry notes and to the Forensic ethics, consent and legal procedure notes. Where it turns on how a prevalence was estimated, or on which design would establish it, it belongs to the Research Methods, Biostatistics and Epidemiology notes. What this chapter supplies is the layer between them: who delivers what, to which population, and how you would know whether they are.

40.1 Four questions locate the service variable

Before deciding what a community psychiatry stem is asking, run the same four questions across it. Between them they cover almost every way an examiner can turn an ordinary clinical picture into a problem about a service, and each opens onto a different part of this chapter.

Level	Offer	Access	Denominator
WHICH TIER IS ASKED TO ACT	WHAT THAT TIER ACTUALLY HOLDS	WHO IS NOT ARRIVING	WHAT THE FIGURE COUNTS

Level asks which tier the question is really being put to: the household and the community worker, the primary care facility and its medical officer, the district hospital, or the specialist institution above them. Offer asks what that tier actually holds, in staff, in supply and in training, as against what the reader assumes it holds. Access asks who is not arriving at all, because a service used by the people who were always going to arrive has changed nothing for the people

who were not. Denominator asks what any figure in the stem counts, who issued it, which period it describes and what it is a proportion of. A well-built stem makes one of the four decisive and leaves the other three as scenery.

RULE **The patient is the setting; the service is the question.** If your answer can be written without naming a tier, a provider, a supply, a pathway or a population, you have answered a question about a person while the stem asked about a system. The clinical picture is usually scenery: it tells you what this person needs and almost nothing about who is meant to provide it. The marks sit in naming the level at which the failure sits, saying what that level does and does not provide, and choosing the move that acts there rather than at the bedside.

40.2 The vignette that resolves on the tier that should act

The commonest stem describes a patient stable for a long time who is not stable now, with the prescription unchanged and nothing new in the history except the travelling. She lives in a village; the clinic that knows her is in the district town; the journey takes most of a working day; somebody has to come with her and lose a day of wages; and the medicine she was given runs out before the next appointment falls due. The reading move is to locate the failure on the ladder of care rather than inside the patient, then ask what the tier nearest her actually holds. A relapse produced by a supply that ran out is a distribution problem; one produced by an appointment the household cannot afford to keep is a geography problem; one produced by a family that stopped believing the treatment was working is a continuity problem. None of the three is answered by changing the drug, and the choice of agent, if the stem really wants one, belongs to the Schizophrenia: management, antipsychotics and adverse effects notes.

The answer that earns marks moves the work downwards and keeps the supervision upwards. Follow-up for a person stable on a settled regimen is the kind of task that can sit with a trained general provider working to a structured protocol, while the specialist tier holds the complicated presentations, the treatment failures and the teaching. That redistribution is what **task-shifting** names, and the intervention guide most often examined for it, with the stepped and collaborative arrangements that make it safe, is developed in this chapter's primary care, task-shifting and stepped-care sections. Where distance is the whole of the problem and a trained provider exists but a specialist does not, a scheduled remote consultation moves the scarce person rather than the sick one, and the conditions attached to doing so sit in this chapter's telepsychiatry section. What a candidate should not write is "refer to the psychiatrist" against a stem whose entire point is that the referral is the thing not happening.

PITFALL

Writing "defaulted follow-up" in the notes without ever asking what the visit costs. The entry travels forward as a statement about the patient's motivation, and every later clinician reads it that way. Ask instead how far the journey is, what the fare and the lost wages come to, who has to accompany her, and whether the medicine lasts until the date you have written on the card. A follow-up plan the household cannot afford to keep has been written for the convenience of the clinic, and it will be read as non-adherence for years afterwards.

40.3 The vignette that resolves on what the programme actually provides

A second family holds the patient constant and moves the question to the district. The stem says the district is covered by the national programme, or that a service was sanctioned for it, and then describes somebody who could not get care there: no clinic ran that month, the psychiatric posts are vacant, the medicines on the district list were not in the store, or the outreach that appears in the annual report has never reached that block. The reading move is to separate **sanctioned coverage** from **functional coverage**. A sanction is an administrative event: an approval, a budget line and a set of posts on paper. Function is a service actually delivered to a population over a stated period, which needs the posts filled, the supply arriving, the staff trained and the activity recorded. The two are counted by different instruments, and the first is far easier to count, which is why it is the figure that travels.

The answer therefore asks what the sanction bought and which of its parts is missing, because the missing part names the level that has to act. Vacant posts are a recruitment and cadre question for the state; empty shelves are a procurement question; untrained staff are a training question; and a clinic that runs while nobody attends is an access question, which is a different vignette again. Naming the missing input is most of the answer, and it is also the honest thing to say about coverage figures in general: what gets published is what the system can count. The components a sanction is meant to buy are set out in this chapter's district-programme section, and the route the money takes to reach them in its financing section.

TRAP **Reading a district count as a service count.** A total of districts approved for a programme is a statement about administrative decisions, not about clinics that opened, and it is quoted as though it described a population served. The same error runs through every coverage figure here: facilities equipped is not facilities functioning, posts sanctioned is not posts filled, and a programme present in a district is not a programme present in every block of it. When a stem offers a coverage figure, the examinable move is to say what it counts before saying what it shows.

40.4 The vignette that resolves on what the figure is a count of

A national number inside a stem is an invitation to interrogate it, not a fact to reason from. The stem quotes something the candidate has heard in a seminar: the proportion of people with a mental disorder who are not in treatment, the calls a national helpline has answered, the density of psychiatrists in the population, the share of the health budget reaching mental health. Then it asks for a conclusion. The reading move is to attach four fields to the figure before using it: who issued it, which period it describes, which population it was measured in, and what it is a proportion of. A figure carrying all four is evidence. A figure carrying none is a rumour with a decimal point in it, and any conclusion drawn from it inherits that weakness.

Two of the traps are structural rather than careless. A cumulative count of contacts with a service, such as the calls a helpline has answered since it opened, counts events and not people: one person in distress may call many times, and no arithmetic on the published total recovers how many individuals it stands for, or how many of them reached care afterwards. And a count with no **valid denominator** cannot be turned into a rate, however inviting the comparison. A national

tally of deaths within an occupational group has no trustworthy count of the people in that occupation to sit beneath it, so a rate built from the two puts an uncertain numerator over an unknown one and reads as precision. The honest sentence names what was counted, says what the denominator would have to be, and stops there. Ask also when the measurement was made, because a household survey describes the period it sampled and not the year in which you are reading it. Where the stem is really asking how such a quantity should be measured, or which design would establish it, that reasoning belongs to the Research Methods, Biostatistics and Epidemiology notes.

The third trap is the one that puts contradictory figures into print. One survey can report the proportion untreated among people with common mental disorders, a second proportion among people with severe mental illness, and a further breakdown for one named condition, and those values are not interchangeable because the populations are not. Each carries a different consequence for which tier should act, since a gap concentrated in severe illness points at the district and specialist tiers while a gap concentrated in common disorders points at primary care. The figures themselves, with the survey that issued them and the populations they describe, sit in this chapter's survey and treatment-gap sections; the disorder-by-disorder breakdown and its epidemiological meaning belong to the Research Methods, Biostatistics and Epidemiology notes.

TRAP **Quoting a survey's figure against a population it did not measure.** One national survey yields several treatment-gap values, one for each population it defined, and the commonest error here is to carry away a single remembered number and attach it to whichever disorder a stem names. The number is then correctly quoted and false at once, which is the hardest kind of error to see, because nothing about the sentence looks wrong. Carry the population with the figure, or carry neither.

40.5 The vignette that resolves on where the pathway to care broke

A third family gives a long history before the first psychiatric contact. A young man has been unwell for years; the family took him first to a healer at a shrine, then to a general practitioner who gave him something for sleep, then to a physician who investigated him thoroughly, and only after all of that to a psychiatric service. The stem is not asking whether the healer did harm, and it is not asking for a formulation of the family's explanatory model at the bedside: that assessment, with the idioms in which distress is described, belongs to the Consultation-liaison, geriatric and transcultural psychiatry notes. It is asking where the **pathway to care** could have been shortened, and by whom.

Read the pathway as a sequence of contacts, each of them a person who could have redirected it. The answer names the point at which recognition failed, the provider standing at that point, and the intervention that acts there: awareness in the community, so that a family reads the problem as one a health service treats; recognition training for whoever the population meets first; a working relationship with the practitioners outside the health system whom a population consults before it consults a clinician; and a service near enough for a referral to be kept. Stigma belongs in this vignette as a service fact rather than an attitude, because it decides who presents,

how late, to whom, and under what description. The strategies, the models of engagement and the evidence for each are set out in this chapter's stigma, literacy and traditional-healing sections.

GOOD TO REMEMBER

Ask every new patient and family two questions before the history proper: who did you see first, and what were you told. The answers give you the route this population actually travels, the contact at which somebody could have redirected it, and the explanation the family is still carrying into your consultation. Collected across a clinic year, those two answers are the closest thing a district service has to a map of its own catchment, and they cost nothing but the asking.

40.6 The vignette that resolves on the population rather than the person

A fourth family describes a death, or a cluster of deaths, and the temptation is to answer with an assessment. A farmer in a district of failed rains and heavy debt has died by suicide using a pesticide kept in the house, and the health administration asks what the service should do. An individual risk assessment is a correct answer to a different question, and the evaluation of risk in the person in front of you belongs to the Somatic-symptom, sexual, impulse-control disorders and suicide notes. What this stem asks is which population lever its facts point at, and they point at several at once: a means kept in every house in the village, an economic shock the health system did not cause and cannot undo, the people the man met before he met any clinician, and a local press that will report the death tomorrow.

The reading move is to separate a **population-level intervention** from a clinical one by asking who receives it. A clinical intervention is delivered to a person already identified; a population intervention is delivered to everyone in a defined group whether or not any individual in it has been identified, which is exactly why it can reach people no service has ever seen. Restricting access to a lethal means, preparing the non-clinicians a distressed person actually meets, agreeing how a death is reported, and acting on the debt and distress that brought the person to the act are all interventions of the second kind. Which one the stem wants is decided by the fact it planted, and this chapter's suicide prevention sections carry what each involves and what the evidence for it is; a vignette answer names the lever and the tier that pulls it, and does not claim a size of effect for it. The same move works on an event that strikes a whole community at once, where an answer built on assessing everybody present individually has answered at a level the event does not permit, and the staged population response sits in this chapter's disaster section.

40.7 The vignette that resolves on what receives the patient

The last family reverses the direction. A man has spent many years in a long-stay institution, his symptoms are settled, the team records that he is fit for discharge, and no family can be traced. Answering with a discharge prescription answers a clinical question that has already been settled; the open question is what receives him. Where people are moved out of institutions faster than the structures they move into are built, the result is **transinstitutionalisation**: the same person held in another closed setting, a custodial one, a shelter or a household with no rehabilitation in it, while the population looks deinstitutionalised because a bed count fell. The rationale, the

outcomes and the lessons drawn from that history are developed in this chapter's deinstitutionalisation section.

So the answer inventories the receiving structures by function: somewhere to live with the support this person needs, something to do in the day, a route towards work where that is realistic, a clinical service that will follow him up without a journey he cannot make, and an entitlement that does not depend on relatives nobody can find. Each is a different service with a different provider, and the provider is often not the state; this chapter's section on Indian models of care sets out the part played by non-governmental organisations, by families and by indigenous resources, and a good answer names who would actually have to be asked. The certification and the statutory entitlement that follow a disability assessment are not this chapter's to state and belong to the Mental health legislation and forensic psychiatry notes.

40.8 Look-alike stems and the answering spine

Consolidation is easiest in pairs. The table sets out the stems most often confused in this territory, against the fact that separates them and the axis the reasoning lands on.

TWO STEMS THAT LOOK ALIKE	THE FACT THAT SEPARATES THEM	THE AXIS IT LANDS ON
A patient relapsing on an unchanged prescription, against one relapsing because the supply ran out	Whether the stem describes a change in the illness or a break in delivery	Level
A district counted as covered, against a district where a clinic actually runs	Whether the figure records an approval or a service delivered over a stated period	Offer
A helpline's total of calls, against the number of people it has reached	Whether the count is of events or of individuals	Denominator
A gap figure for common mental disorders, against one for severe mental illness	Which population the survey measured before any proportion was calculated	Denominator
A long delay before first psychiatric contact, against a family that refused care	Whether anyone met earlier in the sequence recognised the problem and could redirect it	Access
A death that asks for a risk assessment, against one that asks for a population response	Whether the stem names an identified person or a defined group	Level

MNEMONIC

Level, Offer, Access, Denominator

The four service variables spell LOAD, which is what a population puts on a system and what a system is built to carry. **L**evel, the tier the question is being put to. **O**ffer, what that tier actually holds in staff, supply and training. **A**ccess, who is not arriving, and why. **D**enominator, what a figure counts, who issued it, when, and out of what. Each opens a different part of this chapter, and a stem is usually built on exactly one.

With the axis identified, an answer that earns marks follows five moves in order.

- Name the level at which the failure sits, before naming any intervention.
- Say what that level actually provides, and which of its inputs is missing.
- Name who is not reaching the service, and at which contact the pathway broke.
- Attach the issuing body, the period, the population and the denominator to any figure you use, or say plainly that the figure does not carry them.
- Say what you would not do, and why the obvious clinical answer leaves the service variable untouched.

That last move is where public health becomes visible to an examiner. A candidate who answers that the patient should be restarted on medication and reviewed monthly has answered a question about one person, and the answer would read the same in any country and any decade. One who answers that the follow-up belongs at the tier the patient can reach, that the district's problem is a vacant post rather than an absent programme, that the figure quoted describes a different population, and that the people worth worrying about are the ones who never arrived, has answered a question about a service. Every vignette in this territory rewards the second answer.

Find the service variable - level, offer, access or denominator - and answer with the tier that has to act and the population it has to reach, never with the next step in a management plan.

41 · Consolidation and quick review

Nothing in this chapter is about a single consultation. Every section here starts with a population carrying more illness than the service in front of it is reaching, and asks what to do about the difference: a national programme built for districts that had no psychiatric service at all, a survey that counted illness the clinics never saw, a suicide prevention lever aimed at people who will never present, a primary care worker asked to do what no specialist is available to do, and a school, a prison or a camp holding people no clinic will ever see. The individual encounter belongs elsewhere. What belongs here is everybody who never reaches one.

Call that difference the **need-to-service distance**: the gap between what a population needs and what its service actually delivers to it. This closing section teaches no new programme and no new figure, because each is taught in the section that owns it. What remains is to state the

argument those sections were arranged to make, and to convert it into something producible from a blank page with the book shut. The argument is that community psychiatry is the management of that distance, and that the reasoning is the same every time: measure the distance honestly, name who is responsible for closing it, choose the intervention that reaches the people who are not reaching care, and find out whether it did. Questions here are set in exactly that shape, which is why an answer built as a list of programme names scores so poorly against one built as a plan.

41.1 One argument, tested against every population here

Read as a list of territories this material is enormous and disconnected: programmes and policy, surveys and burden, suicide prevention, stigma, rehabilitation, integration into primary care, disasters and special settings, prevention and financing and digital services, culture and traditional healing. Read as one argument it is a single claim tested across all of them. The claim is that each section describes a distance between a population and a service, that the distances differ only in what produced them, and that the same four moves answer every one.

They recur because there are only so many ways a person with a treatable illness fails to end up in treatment. The illness is not recognised as illness. It is recognised, and reaching help is shaming, or is better answered somewhere the family already trusts. It is recognised and acceptable, and there is no service within reach or nobody in it trained to treat. It is affordable to nobody in the household once travel and lost earnings are counted. Contact is made once and never repeated, so what was received was never adequate to the illness. And in some populations the person sits somewhere no service goes at all. Every section here is one of those situations wearing a programme's name, which is why the reasoning does not change as you move through them.

RULE **A service is adequate or inadequate only against a stated population, never on its own description.** The service's own account tells you what it offers; the population tells you what share of the need that offer reaches. An answer that lists a programme's components and stops has described the offer and said nothing about the distance, which is the only thing the programme exists to change.

41.2 Need side or delivery side: the question that opens every answer

Before anything is planned or measured, the distance has to be sorted into one of two kinds, because the two are closed differently and are confused constantly. A **need-side distance** means the population is not arriving: the illness is not recognised, help-seeking is unacceptable, the first person consulted is not a mental health service, or coming costs more than the household will pay. A **delivery-side distance** means the population arrives and the service cannot deliver: no trained worker, no medicine on the shelf, no follow-up, no route onward for the person who does not improve. Both are present in almost every district and neither excludes the other, but at any one moment one of them is binding, and effort spent on the other moves nothing.

The distinction decides the intervention. A need-side distance is answered on the population's own ground: awareness and literacy work, contact with people who have been unwell, respectful arrangements with the practitioners and healers a family consults first, removing cost at the point

of use, and taking the service to the school, the workplace, the prison or the camp. A delivery-side distance is answered inside the health system: training and supervising the workers already posted, moving a bounded task to a worker who can safely do it, securing medicine supply, opening a remote link to a specialist who cannot be posted locally, and arranging the route by which somebody who is not improving is escalated. Candidates who cannot separate the two propose more specialists for a population that was never going to come, and awareness work for a population already waiting outside a clinic with an empty pharmacy.

It also decides what to measure. A population survey measures need and says nothing about what any service did; a service record measures what the service did and says nothing about the people who never came. Only a figure with a service numerator over a population denominator measures the distance itself, which is why **coverage** is the most useful and least available number in this whole area.

41.3 The four moves, and why they run in that order

Once the distance is named and sorted, one sequence organises the answer whichever section it came from. It is a revision scaffold and a planning discipline rather than an official procedure, and it yields in one situation: where a population is in immediate danger, act first and measure afterwards, which is the standing position in the first response to a disaster. Outside that the order matters, because each move is close to meaningless unless the one before it has been made.

- Measure the distance, and measure it honestly. State who issued the figure, what period it refers to, which population it describes, and what it is a proportion of. A figure missing any of the four is not a measurement, it is a quotation. Where it does not exist, say so, and name what would have to be counted to produce it.
- Name who is responsible for closing it. Every distance here sits at a level that can act: a household, a village-level worker, a primary care team, a district unit, a state, a national programme, an employer, a school, an insurer, a court. An intervention proposed without a level that owns it is a suggestion, and it does not survive the meeting it is proposed in.
- Choose the intervention that reaches the people who are not reaching care. Interventions differ far more in whom they touch than in how well they work in those they touch, so selecting on effect alone selects, reliably, whatever works beautifully in the people already coming.
- Find out whether it worked, and design that before it starts. What was the distance before, what is it now, in the same population and against the same denominator. Deciding this at the end leaves only whatever the programme happened to record, which is almost always what it did rather than whom it reached.

Reversing any two produces a failure recognisable in a real service. Choosing the intervention before measuring the distance is how a district acquires the service its planners find most familiar rather than the one its population lacks. Choosing it before naming the owner produces the pilot that works and never scales, because nobody whose job it was ever ran it. And measuring, naming and delivering without ever evaluating is the commonest of the three: it is how a programme

accumulates a shelf of **activity returns** and cannot answer, when finally asked, whether the distance it was built to close has moved.

Distance	Duty	Door	Difference
MEASURED WITH ITS OWN DENOMINATOR	THE LEVEL THAT OWNS CLOSING IT	THE ROUTE THAT REACHES THE ABSENT	THE SAME MEASURE, TAKEN AGAIN

41.4 A frame that names every distance

Because the distances recur and the four moves do not change, one retrieval frame serves the whole chapter. Hold it as a set of prompts rather than as a set of topics: under pressure the failure is rarely that a candidate cannot name a barrier, it is that they answer on the barrier the question made obvious and never scan for the others.

MNEMONIC

Distance, Duty, Door, Difference

The four moves of every answer in this area, in the only order that works. **Distance**: measure the gap between what the population needs and what the service delivers, carrying the body that issued the figure, the period it refers to, the population it describes and the denominator underneath it, and declare it missing when it is missing. **Duty**: name the level that owns closing it, because an intervention with no owner has no route to delivery. **Door**: choose the intervention by whom it reaches, not by how well it works in those already attending. **Difference**: measure the same distance again, in the same population and against the same denominator, and treat a count of what the service did as no answer at all. The same four, asked in the same order, of a national programme, a district unit, a school, a prison or a helpline.

Each distance is taught in full in the section that owns it, with its own evidence and its own figures. The frame below is a retrieval index rather than a teaching table, and it gives every row the same three columns deliberately, because the argument is that the columns do not change.

THE DISTANCE	WHICH SIDE IT SITS ON	WHAT CLOSES IT, AND WHAT WOULD SHOW IT CLOSED
The illness is not recognised as illness	Need side, before any decision to seek help	Awareness and literacy work, and recognition training for whoever is consulted first. Shown by a change in who arrives, not by a count of people who attended a talk
Help is not sought, or is sought elsewhere first	Need side, and decisive where it operates	Contact-based work, change to the rules that disadvantage, and working arrangements with the healers seen before any psychiatrist. Shown by who arrives and how late, not by an attitude score
Care is unaffordable once fee, travel and lost earnings are added	Need side, and the one most often left out of a plan	Cost removed at the point of use, financing that follows the patient, care delivered nearer to people. Shown by uptake among the poorest, not by a scheme existing
No service within reach, or nobody in it trained to treat	Delivery side	Integration into general and primary care, task-sharing with supervision, secure medicine supply, remote links to a specialist. Shown by people treated, not by workers trained
Contact is made once and never repeated	Delivery side, and the most expensive to repair	Follow-up that is somebody's named job, a register of who is due, and a route that escalates when nothing improves. Shown by retention, not by first attendances
The person sits where no service goes: institution, prison, camp, shift, classroom	Both sides at once, which is why these populations are reached last	A service designed for the setting, with the setting's own authority attached. Shown by the population inside the setting, not by the district total

41.5 The measurement layer, and what each figure is for

Numbers are this chapter's connective tissue and its commonest source of error, and the error is almost never arithmetic. Each figure is best held as a **measurement proxy**: a countable thing standing in for something that matters and cannot be counted directly. What a figure is a proxy for tells you what a healthy-looking value fails to prove, which is what is being tested whenever a question hands you a national number.

WHAT IS COUNTED	WHAT IT IS A PROXY FOR	WHAT A GOOD VALUE FAILS TO PROVE
Prevalence in a population survey	How much illness exists, in the population and period surveyed	Anything at all about services. A survey counts people, whether or not one of them ever reached care
Units sanctioned, approved or established	An administrative decision that a service ought to exist	That any of them opened, was staffed or saw a patient. A sanctioned unit is an intention, not a service
Activity: clinics held, workers trained, calls answered	Effort expended, which is the easiest thing any programme can count	That anybody was reached who was not already being reached. Contacts are not people, and one person contacting twice is counted twice
Coverage: those treated over those with the condition	The distance itself, and the only figure measuring it directly	That treatment was adequate. Who counts as treated depends entirely on the adequacy rule applied
Outcome in those treated, such as recovery or remission	Whether the treatment works in the people who received it	That coverage moved. An intervention can produce excellent outcomes in everyone it reaches and leave the distance where it was
Money allocated, released or spent	Stated priority, and separately the capacity to convert money into service	That the service exists. These are three different figures, and only the third attaches to something that happened

Held this way the numbers stop being a list to memorise and become a set of questions, and the drill worth doing is the reverse one. Take a figure of the kind a question will hand you and, before commenting on its size, produce its population, its period, its issuing body and its denominator. Where one of those cannot be supplied, naming the missing field is a stronger answer than a confident number.

41.6 When closing the distance is not the aim

Two situations sit outside the argument as stated, and applying the frame mechanically to either is wrong in a way that is easy to hear. The first is the **population lever**, the intervention that never touches care-seeking at all. Restricting access to a lethal method, changing how a death is reported, designing work so that it does less harm, amending a rule that disadvantages people with mental illness: none of these depends on anybody recognising an illness, deciding to seek help or reaching a service. They act on everybody at once, they are usually delivered by somebody outside health, and their effect appears in no service record. The question to ask of them is not who was reached but whether the exposure changed, and the level responsible is often a ministry, a regulator or an employer rather than a health service.

The second is the distance that cannot be measured with the counts available. Some of the most quoted figures in this field have no valid denominator: a cumulative total of contacts with a service is not a count of people, and a count of deaths in an occupational group has no matched population figure to sit over. The honest position is that the numerator is real, the denominator does not exist in matchable form, and no rate may be built from the pair. Saying so is not a failure to answer. Building the rate anyway is the failure, and it is the one that gets quoted onward by everybody who hears it.

41.7 The two ways candidates lose these marks

Two failures account for most of what is lost here. The first is **catalogue answering**: reciting programmes, components and structures without once saying what distance any of them addresses or how anyone would know it had closed. It reads as prepared, it fills a page, and it answers an administrative question nobody asked. The second is quoting a national figure without the population it belongs to, which is worse than it looks, because a correctly transcribed number attached to the wrong group is a false statement and no arithmetic check will catch it.

TRAP **Answering a service question with more specialists.** Asked how to improve care for a district, a population or a special setting, candidates reach almost reflexively for more psychiatrists and more clinics, which is a delivery-side answer given without checking which side the distance is on. Where the population is not arriving, capacity added at the clinic door changes nothing and shows up in no measure. An answer that names the side first, then the level responsible, then the intervention, then the measure that would show it worked, has earned most of what is available whatever it eventually proposes.

PITFALL

Reading a busy service as a reached population. A unit that is full, whose staff are exhausted and whose returns look excellent, tells you what happens to the people who arrived and nothing whatever about those who did not. The two move independently, and the busiest services are often the ones with the least idea of their own catchment. The habit worth keeping is to ask, of any service you work in, what share of the people in its area with the condition it treats has ever reached it. Where nobody can answer, that is the finding.

41.8 Where this chapter stops, and what to carry in

An answer is judged as much by what it declines to teach as by what it covers. Consolidating this material means consolidating its edges, because several neighbouring chapters treat the same programmes, populations and deaths from a different angle. Each of the following is worth rehearsing as a handover: named where it touches the question, passed to its owner in a sentence, and left there.

- The provisions of mental health legislation, the advance directive, the nominated representative, the review machinery and the decriminalisation of attempted suicide belong to the Mental health legislation and forensic psychiatry notes. What is kept here is the service consequence: who arrives, what the service records, and what capacity the change assumes.

- Professional conduct, consent, capacity at the bedside, custodial and court procedure and the statutory arrangements for disaster response belong to the Forensic ethics, consent and legal procedure notes.
- Incidence, prevalence, study design, sampling and the derivation of any rate belong to the Research Methods, Biostatistics and Epidemiology notes. This chapter uses those measures as instruments and never derives them.
- The assessment and management of suicide risk in the person in front of you belongs to the Somatic-symptom, sexual, impulse-control disorders and suicide notes. What is kept here is the population lever.
- Cultural formulation at the bedside, idioms of distress, the named culture-bound presentations and the assessment of somebody who has consulted a faith healer belong to the Consultation-liaison, geriatric and transcultural psychiatry notes. What is kept here is the pathway and the partnership.

GOOD TO REMEMBER

Before writing a line of any service proposal, put the four moves down as four headings and fill them in that order. The discipline shows in the fourth: deciding what measurement would demonstrate a change usually forces a change in the intervention itself, because attractive proposals often turn out to move nothing anybody can count. Choosing the measure last, once the plan is fixed, is how a service ends up with years of returns and no evidence.

Carried into the examination, all of it collapses to a sequence that survives when the detail fails. Name the population and the distance. Say which side is binding. Name the level responsible for closing it. Choose the intervention on whom it reaches. State the measure that would show the distance had moved, and say plainly where that measure does not exist. Consolidation buys not more facts but faster access to the ones already learned, so that under pressure the answer arrives as a method rather than as a half-remembered list of programme names.

Community psychiatry is the management of the distance between what a population needs and what its service delivers: measure the distance with its denominator, name the level responsible for closing it, choose the intervention by whom it reaches rather than by how well it works in those already attending, and then measure the same distance again.

PART XI

XI

Medical and Systemic-Disease Psychiatry

Psychiatry at the medical interface: the heart, the lungs, the failing organ, the transplant and the cancer, and the discipline of routing what this chapter does not own.

01 Cardiac, respiratory and critical-illness psychiatry

02 Renal, transplant and perioperative psychiatry

03 Oncology, palliative and paraneoplastic psychiatry

04 Autoimmune and inflammatory systemic disease, and burn injury

05 Consolidate and apply

VISUAL 11.1



The organ interface changes psychiatric decisions

- Disease physiology
- Psychiatric symptoms
- Psychotropic exposure
- Organ adjustment
- Liaison formulation
- Monitoring

Medical psychiatry asks how organ disease, psychiatric symptoms and treatment exposure alter one another and what must be monitored.

Source lock: day-40.final.html#sec-t01

VISUAL 11.2



Critical illness has a long psychiatric shadow

- Delirium
- Anxiety and PTSD
- Physical recovery
- Cognition
- Family experience
- Follow-up

Post-intensive-care outcomes include psychological, cognitive and physical problems in patients and families and require longitudinal follow-up.

Source lock: day-40.final.html#sec-t03

VISUAL 11.3



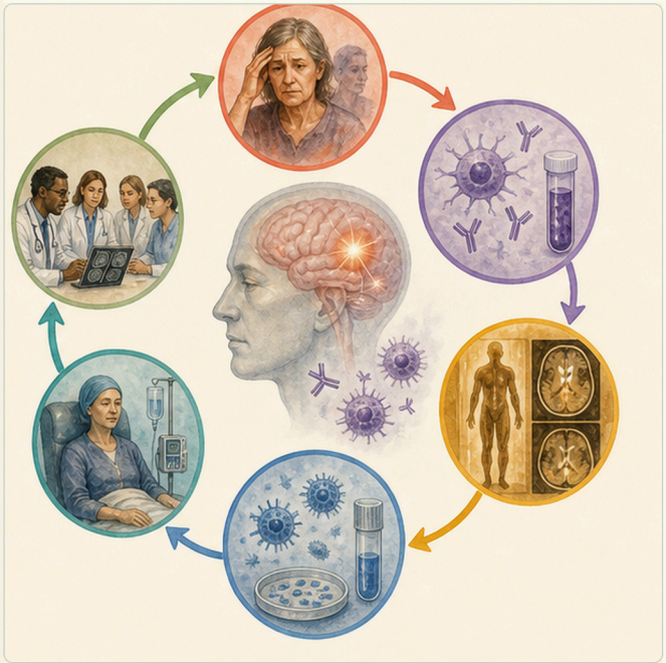
Palliative psychiatry separates overlapping distress

- Depression
- Demoralisation
- Delirium
- Existential distress
- Symptom care
- Shared goals

Palliative psychiatry distinguishes depression, demoralisation, delirium and existential distress while aligning symptom care with the patient's goals.

Source lock: day-40.final.html#sec-t06

VISUAL 11.4



Systemic disease can signal brain inflammation

- Neuropsychiatric red flags
- Immune evidence
- Malignancy search
- Exclude infection
- Immunotherapy
- Coordinate specialties

Autoimmune, inflammatory and paraneoplastic syndromes are approached through red flags, staged exclusion, immune evidence and coordinated treatment.

Source lock: day-40.final.html#sec-t16

PAPER IV

Medical and systemic-disease psychiatry

Four bands . one complete notes document

Cardiac, respiratory and critical-illness psychiatry

1 · The organ interface: what this chapter owns, and what it deliberately points to

At the interface with medicine the psychiatrist is rarely asked what the diagnosis is, and is almost always asked which of three things is producing the picture. The medical illness may be generating the psychiatric symptoms directly. The treatment given for that illness may be generating them. Or the patient may have a psychiatric disorder of their own, arriving in a body that happens to be unwell. Those three answers oblige three different actions, and this chapter is built to make the choice between them quicker and more defensible. This orientation installs that **triage**, adds the fourth question the triage forces, which is who teaches a construct once it has been placed, and draws the map of what follows.

One feature of this chapter should be stated at the outset rather than discovered halfway through. It is deliberately short for the number of constructs it names, because several of them are named here and taught elsewhere in the series. That is a design decision with a defence, and the defence is set out below in full, because a reader who does not understand why a chapter points will read the pointing as a shortfall.

1.1 The three-way question, and what each answer obliges you to do

The first limb is the illness acting on the brain. A systemic disease can produce psychiatric symptoms through mechanisms that require no psychological step at all: inflammation reaching the central nervous system, an immune response raised against a tumour that also recognises neural tissue, hypoxia and carbon dioxide retention, the accumulated products of failing kidneys, endocrine derangement, small-vessel injury, a surge of circulating catecholamines. When this limb is the answer, the psychiatric task is largely to establish it and to hand the patient back for treatment of the disease, because no psychiatric intervention will hold while the process continues.

The second limb is the treatment acting on the brain. Modern medical treatment is potent and it is neuroactive, and the psychiatric consequences of it are not rare accidents but predictable properties of certain agents and certain procedures. Immunosuppression, cytotoxic treatment, interferon, corticosteroids, anaesthesia and surgery itself, oxygen and ventilatory support and implanted devices all change the mental state, each with its own syndrome and its own interval. When this limb is the answer, the decisive evidence usually sits in the drug chart and the procedure record rather than in the mental state examination, and the action is a conversation with the treating team about the exposure.

The third limb is a psychiatric disorder in its own right, occurring in a person who is medically unwell. This limb has two internal shapes that behave alike for management purposes. The

disorder may be one the patient would have developed anyway, whose timing merely coincides with the admission. Or it may be a response to the illness, to what the illness has taken, and to what it will take next. Both need treating as psychiatric conditions, and both are vulnerable to the same argument, which is that a person in this situation would reasonably feel that way. When this limb is the answer, the psychiatric task is the largest and the least likely to be delegated.

Two properties of the triage matter more than the categories themselves. The limbs are not exclusive, and in a medically unwell patient more than one is usually operating at once, so the working question is not which limb is true but which limb is doing the most damage now and which, if missed, is dangerous. And the limbs are answered by different evidence: the first by the disease course and its physiology, the second by exposure and timing, the third by the patient's history, baseline and account of themselves.

Illness	Treatment	Independent disorder	Owner
READ FROM THE DISEASE COURSE	READ FROM EXPOSURE AND TIMING	READ FROM BASELINE AND HISTORY	WHO TEACHES WHAT YOU HAND ON

RULE **State which limb you are in, and say what evidence put you there.** An answer at this interface that names a psychiatric syndrome without placing it against illness, treatment or independent disorder has skipped the step the chapter examines. The examiner's follow-up is almost always the same question in different words: what else could be producing this, and what would separate them?

1.2 The examination technique the triage runs on

The triage is decided by history far more often than by investigation, and by a small number of manoeuvres that are cheap and frequently skipped. Learning them as a fixed order is worth more than learning any single association in this chapter, because the order transfers to every section in it.

Establish the baseline first, from someone other than the patient. What was this person like a month ago, at work, at home, in conversation? A medically unwell patient is observed by staff who never met the well version, so the comparison that would make the change visible does not exist on the ward unless a psychiatrist goes and gets it. Without a baseline, a description of the present state is the only thing available, and a present state cannot establish a change.

Then establish **tempo**. Every psychiatric illness has a characteristic time signature, and the useful reading is the mismatch: a presentation moving faster than the condition it resembles, or continuing to move while adequate treatment is being given, or following a slope rather than an episode. Tempo costs nothing but a careful timeline with corroboration, and it is the reading most often left vague.

Then lay two timelines side by side. The first is the medical course, with its events dated: the acute episode, the operation, the transfer, the deterioration. The second is the drug chart, read for start dates rather than for the current list. Laying the psychiatric change against those two timelines is the single highest-yield act at this interface, and it is what separates the first limb

from the second. A change that begins with the illness and tracks its severity points one way. A change that begins after an agent was introduced, and in the interval characteristic of that agent, points the other.

Ask what was stopped as well as what was started. An admission removes access to alcohol, to nicotine and to whatever else the patient used at home, at a predictable hour, and it frequently interrupts the patient's own long-term psychotropic treatment on the ward round of a team that does not prescribe it. A withdrawal syndrome or an abrupt discontinuation, declaring itself on a medical ward, is read by people looking for a medical explanation as a complication of the illness.

Attend to consciousness and attention before anything else in the mental state. A fluctuating disturbance of attention and awareness changes the entire answer and moves the problem into territory this chapter does not teach, and the whole of that assessment, including any bedside instrument, rating scale or drug choice for it, belongs to the stroke, traumatic brain injury, delirium and focal brain syndromes notes. Recognising the shape and handing it on is the examinable act here.

TRAP **Settling for understandable.** The commonest error at this interface is not a wrong diagnosis but an abandoned one: the low mood is judged a reasonable response to a serious illness, and the assessment stops. Comprehensibility is not a diagnostic exclusion and would not be accepted for any other comorbidity. It costs more than usual here, because the psychiatric state is often what removes the patient from the treatment, the rehabilitation and the adherence that would change the medical trajectory.

1.3 Why a chapter that points beats a chapter that teaches everything thinly

The alternative to routing a topic away is not depth, it is a thin second copy, and a thin copy is worse than an absence because it does not feel like one. A paragraph read on a topic taught properly elsewhere leaves the reader recognising every term, able to defend none of them, and unable to tell which of their own facts are the thin ones. A half-learned fact sits in memory looking exactly like a fact learned well. It survives the first question and fails the follow-up, which is where the marks are, and it fails at the bedside for the same reason.

An explicit omission behaves differently. A reader who has been told that a construct is taught in a named chapter knows precisely what they do not yet have, and can act on it. A reader who met a thin paragraph knows nothing of the kind. That difference is also what makes routing checkable rather than merely tidy: whether a chapter mentioned a topic is a weak test that a chapter can pass while leaving its reader unable to answer anything, whereas whether the reader can name the owner is a test with a right answer.

There is a second argument from the reader's side, and in the last weeks before an examination it is the stronger one. Revision time is fixed, and a chapter that duplicates an adjacent chapter spends that time twice on one topic and not at all on the topics only it can teach. The material this chapter alone carries, which is the systemic and inflammatory disease that presents psychiatrically before anyone knows it is present, is exactly the material that goes unrevised when a chapter pads itself out to look complete. Shortness here buys depth where depth is scarce.

A third argument concerns the text rather than the reader, and this chapter's routing section makes it in full: two accounts of one topic, written by different hands, drift apart, and one owner with one account is the only arrangement in which two pages of the same book cannot disagree.

The fourth argument is the one that matters clinically, and it is why the routing question is not bookkeeping. Liaison work is **routing**. A referral arrives from another team, written by someone unsure what to ask, and the psychiatrist's first act is to decide whose question it is: the treating team's, the psychiatrist's, or a third specialty's altogether. A resident who has practised deciding where a construct belongs has been practising the actual skill, and one who has memorised a paragraph about each has not. The map that follows is therefore content, not furniture.

RULE **A declared gap is a correct answer, and a supplied specific is not.** Where this chapter does not hold a figure, a threshold or a management sequence, it says so and names the kind of source such a statement would rest on. In an examination the equivalent move is to state the direction confidently, state that the magnitude is not something you can quote, and go on. That answer is stronger than a number produced from recollection, which is unverifiable and frequently wrong.

The map of routed constructs is drawn once, on the plate below, and the routing section that follows gives each one an exact address and a single boundary sentence stating what this chapter withholds about it.

1.4 The map of the chapter, and the question each section answers

The chapter is organised into four bands. The first covers cardiac, respiratory and critical-illness psychiatry. The second covers renal, transplant and perioperative psychiatry. The third covers oncology, palliative and paraneoplastic psychiatry. The fourth covers autoimmune, rheumatological and inflammatory systemic disease. A section carrying the psychiatric care of the severely burned patient sits alongside them, and a routing section carries the constructs taught elsewhere. Each section below is named by the question it answers, because a question is easier to hold in mind than a title and it is what an examiner is actually setting.

SECTION	THE QUESTION IT ANSWERS	LIMB OF THE TRIAGE IT MOSTLY TRAINS
Heart failure, arrhythmia and cardiac devices	Which of the somatic symptoms belong to the failing circulation, how do rhythm presentations separate from panic, and what does an implanted device do to a person and a household?	Illness, and independent disorder
Obstructive airways disease, asthma and chronic respiratory failure	Is this breathlessness or is it panic, and which behavioural levers change the trajectory?	Illness, and independent disorder
The stress-heart interface	Can an acute psychological event produce a structural cardiac lesion, and why is the presentation mistaken for a coronary event?	The arrow reversed, running from the psychological event to the organ
Transplant psychiatry and immunosuppressant neurotoxicity	What must a pre-transplant psychiatric opinion contain, and which new neuropsychiatric states belong to the drugs that keep the graft?	Treatment
Perioperative psychiatry and postoperative neurocognition	In what order do the perioperative questions arrive, and which cognitive construct is being named after the operation?	Treatment
Cancer-drug neuropsychiatry	Which agent, which syndrome, and over what interval does that syndrome appear and get looked for?	Treatment
Paraneoplastic neuropsychiatric syndromes	When should a psychiatric presentation start a directed search for a malignancy nobody has found yet?	Illness, before the illness is known
Neuropsychiatric lupus	Does the syndrome in front of you belong to the disease at all, and how is that attribution reasoned when no test settles it?	Illness, with attribution as the whole difficulty
Neurosarcoidosis, cerebral vasculitis and antiphospholipid syndrome	How do you answer on a rare systemic cause without inventing a phenotype for it?	Illness, and evidential discipline
Coeliac disease, inflammatory bowel disease and the gut-brain axis	What do the association and bidirectionality findings actually state, and what will their design bear?	Illness, and independent disorder

SECTION	THE QUESTION IT ANSWERS	LIMB OF THE TRIAGE IT MOSTLY TRAINS
Burns psychiatry	What is the psychiatric task at each stage of the burn timeline, and what question does the injury itself raise about what preceded it?	Independent disorder, and response to injury
Organ-interface routing	Who teaches the constructs this chapter names and does not teach, and what exactly is withheld here?	The fourth question

Read the third column across and the lens becomes visible in the shape of the chapter. The sections on paraneoplastic syndromes, lupus and the rare systemic autoimmune causes are extended training in the first limb, and specifically in its hardest form, where the illness has not yet been identified and the psychiatrist is the person who must raise the possibility. The sections on transplantation, the perioperative period and cancer drugs are training in the second limb, where the answer is in the exposure and the interval rather than in the mental state. The cardiac, respiratory, bowel-disease and burns sections spend most of their length on the third limb, on psychiatric disorder in a person who is genuinely and often permanently unwell, and on the reasoning that keeps it treatable rather than merely understandable.

One section sits outside the scheme on purpose. The stress-heart interface runs the arrow the other way, from an acute psychological event to a demonstrable cardiac lesion, and it is in the chapter because it is the clearest available demonstration that the traffic at this interface is not one-directional. Holding that section next to the triage is useful precisely because it is the exception.

1.5 What this chapter does not teach, and where each of those lives

Five boundaries are enforced throughout, and knowing them in advance saves a reader from searching this chapter for material it has deliberately excluded.

- Prescribing at the interface, meaning organ-adjusted dosing, interaction rankings and the choice of agent in cardiac, hepatic, renal or transplant disease, belongs to the psychopharmacology notes. This chapter teaches everything that precedes a prescription and adds no second version of the prescription itself.
- The assessment and management of the acute confusional state belongs to the stroke, traumatic brain injury, delirium and focal brain syndromes notes, including every bedside instrument, rating scale and drug choice for it. Several sections here reach that boundary, at the end of life, after an operation, in critical illness and in the early burn injury, and each marks the hand-off rather than crossing it.
- Capacity belongs to the mental health legislation notes. This chapter says only that capacity is decision-specific and time-specific, which is enough to reason with, and it states no statutory test.

- The generic immune-brain framework, the workup of autoimmune encephalitis and the neuropsychiatry of interferon-beta belong to the neurology and psychiatry interface notes. What this chapter owns on the immune side is the disease-specific presentation and the decision to suspect, not the general model and not the investigative protocol.
- The constructs defined by a clinical setting rather than by an organ, being the coronary care unit, the intensive care unit, the dialysis unit, the oncology ward and the palliative care unit, belong to the consultation-liaison, geriatric and transcultural psychiatry notes, and the psychiatric disturbance that follows a corticosteroid belongs to the endocrine and metabolic psychiatry notes. Those two owners take the constructs this chapter routes, and the routing section names each one with its boundary.

None of this is a shortfall being managed politely. What those constructs share is a referral pathway and a liaison relationship, not a pathophysiology, so they belong with the chapter that teaches liaison work. The one that follows a drug follows the drug. The classification is working, and stating it out loud is the difference between a boundary and an omission.

1.6 How to use this chapter in the weeks before the examination

Read the routing plate before anything else, so that the constructs this chapter will not teach are settled and stop feeling missing. Then read each section for its discriminators rather than for its lists, because every section here is built around a small number of separations that decide the answer, and the lists around them are recognisable at the paper and not defensible under questioning.

Practise the triage out loud on whichever medical patient you happen to be asked about, and practise it as a four-step answer rather than as three categories. The four cues below are a teaching scaffold for that sequence rather than a validated instrument, and the sequence itself is the answer structure this whole chapter installs. It works on a construct you have never met, which is more than can be said for a memorised list.

MNEMONIC

Limb, Evidence, Shift, Owner

Limb: which of the three the picture mostly sits in, the illness acting on the brain, the treatment acting on the brain, or a psychiatric disorder in its own right in a person who is unwell. **Evidence**: what placed it there, which is the disease course and its physiology for the first limb, exposure and timing for the second, and the patient's baseline, history and account of themselves for the third. **Shift**: what would move you to another limb, said out loud, because the limbs are not exclusive and the working question is which one is doing the most damage now. **Owner**: which chapter takes whatever you have just handed on, named rather than rebuilt from memory.

IN INDIA

No India-specific epidemiology is held for this orientation and none has been invented. The Indian point here is one of service structure. In most Indian general hospitals the psychiatrist arrives at the medical bedside on request, usually late, usually to a referral question written by a team under pressure, and often without the previous notes, the collateral history or the drug chart's start dates to hand. Every manoeuvre in the triage above depends on information the referring team holds and the referral letter rarely contains, so the practical skill is asking for it explicitly and early rather than working from the mental state alone. That is also the argument for reading the consultation-liaison material before a liaison posting begins.

GOOD TO REMEMBER

Three questions at the bedside and one afterwards. Is the disease doing this? Is the treatment doing this? Is this a psychiatric disorder in a person who is unwell? And then, before writing: whose chapter answers the part I am handing on?

This chapter teaches a triage and a map, and it teaches them together, because at the medicine interface deciding where a problem belongs is most of deciding what to do about it.

2 · Depression and anxiety after myocardial infarction and in coronary artery disease: prognosis and safe antidepressant choice

Six of the topics in this chapter are addresses rather than lessons, and saying so plainly is the most useful thing this section can do. Depression and anxiety after myocardial infarction and in coronary artery disease, intensive care psychiatry and the psychological sequelae of critical illness, psychiatric assessment in end-stage renal disease and on dialysis, depression and demoralisation in cancer, palliative care psychiatry, and steroid-induced psychiatric disturbance are each named in this chapter and taught in none of it. Each has an owner elsewhere in the series. This section hands the reader that owner, states what is withheld here, and explains why withholding it is honest rather than lazy.

2.1 Why a routed chapter is more honest than a complete-looking one

The alternative to routing is not depth. It is another thin copy. A chapter that gives every adjacent topic a paragraph produces a reader who recognises every term on the paper, can defend none, and cannot tell which of their facts are the thin ones. That is the harm, and it is not obvious: a half-taught fact does not feel half-taught. It sits in memory looking like a fact learned properly, and it fails at the point of use rather than the point of reading.

The second harm is duplication, and it is invisible while writing. Two accounts of one topic, written by different hands against different sources, do not stay identical. They drift, and a reader who has met both cannot tell which is authoritative. One owner and one account is the only arrangement in which two pages of the same book cannot disagree.

The third point is about how completeness gets checked, which is why routing must be declared out loud rather than practised quietly. Asking whether a chapter mentions a topic is an existence check, and a chapter can pass it while leaving the reader unable to answer anything. Asking whether the reader can name the owner is the better test, because it is checkable. A **routed construct** turns an implicit omission into an explicit one, and an explicit omission can be acted on. Silence cannot.

A fourth reason is particular to this chapter. Five of the routed topics are not defined by an organ at all. They are defined by a setting: the coronary care unit, the intensive care unit, the dialysis unit, the oncology ward, the palliative care unit. What they share is a referral pathway and a liaison relationship, not a pathophysiology, so they belong together in the chapter that teaches liaison work. The remaining topic is defined by a drug and follows the drug. Routing here is the classification working correctly, not a shortfall being managed politely.

RULE **Each routed construct gets one exact cross-reference and one boundary sentence here, and nothing else.** This section supplies no incidence, no prognosis, no management sequence and no dose for any of them, and that restriction is its design, not its limitation. A sentence here that supplied one would be the second copy this arrangement exists to prevent.

2.2 The routing plate, and how to read it

The plate below is this section's content, and each row holds three things.

- The construct, worded as the examination will name it.
- The chapter that owns it, which is what the reader is meant to leave with.
- The single sentence stating what this chapter will not say about it, so that the omission is explicit rather than silent.

The rows are addresses with a warning label attached; they will not survive being revised from.

The owning-chapter column is repetitive, because most of these are liaison presentations pointing at the same place. Some rows carry a split, because a construct can have its assessment owned by one chapter and a single limb owned by another, and the plate says so.



Fig 40.1 The routing map, which is the whole visible content of the organ-interface routing section. Six of this chapter's constructs have owners in earlier chapters, and drawing that once is what makes a short chapter honest rather than evasive. Read the shape before the words: five rows converge on a single destination panel, so the reader sees at a glance that consultation-liaison psychiatry, not this chapter, holds post-infarction mood, critical-illness sequelae, dialysis, cancer and palliative care. The sixth row breaks away because steroid psychiatry belongs to the endocrine and metabolic chapter, and that break is the only thing the second colour means. What each row prints in amber is the entire permitted output: a naming sentence, the exact cross-reference, and a full stop. From memory you should be able to say, for any of the six, which chapter owns it and that this one adds no figure, no threshold and no ladder of its own.

2.3 What this section cannot establish, and how to answer without it

Two things a fuller account would supply are declared here as gaps, because a **declared gap** is a correct answer and an invented figure is not.

The first concerns anxiety rather than depression after infarction. A statement of its prognostic effect would rest on a meta-analysis in the cardiology literature, identified for this series but not opened in quotable form. The consequence is stated rather than concealed: this chapter attaches no prognostic figure to post-infarct anxiety, and a candidate should attach none from memory. What can still be said is that anxiety and depression after infarction are separate constructs on separate evidence, and that an answer treating them as interchangeable makes a claim it cannot support.

The second concerns the desire for hastened death in advanced illness. No current guideline recommendation, and no primary intervention study, could be opened for this series with text usable on its own terms on assessing or managing it; such a claim would rest on an oncology

clinical-practice guideline group, or on investigators of primary palliative-care psychotherapy trials. The absence of a quotable recommendation licenses no inaction: a statement of that kind is met as a communication carrying meaning rather than a symptom to be corrected, and decision-making capacity, which is decision-specific and time-specific, is doctrine owned by the Mental health legislation and forensic psychiatry notes. What is withheld here is a protocol, not the obligation to respond.

2.4 The rows in words, with the boundary stated for each

Depression and anxiety after myocardial infarction and in coronary artery disease is owned by the Consultation-liaison, geriatric and transcultural psychiatry notes, and the choice of antidepressant in established cardiac disease is owned by the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes. This chapter states neither the size of the prognostic effect nor the name of any preferred agent.

3 · ICU psychiatry, post-intensive-care syndrome and psychological sequelae of critical illness

Intensive care psychiatry and the psychological sequelae of critical illness are owned by the Consultation-liaison, geriatric and transcultural psychiatry notes, and delirium in the critically ill patient is owned by the Stroke, TBI, delirium and focal brain syndromes notes. The route follows from the shape of the construct, since post-intensive care syndrome encompasses physical, cognitive, and psychological impairments that persist in patients after discharge from an intensive care unit. Only one of those domains is psychiatric, so a chapter organised by organ could own only a slice of it. Its assessment, its course and its treatment appear here in no form.

4 · Psychiatric assessment and management in end-stage renal disease and on dialysis

Psychiatric assessment and management in end-stage renal disease and on dialysis is owned by the Consultation-liaison, geriatric and transcultural psychiatry notes, and dose adjustment in renal impairment is owned by the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes. The routing follows the service design rather than the diagnosis, since psychological support and interventions should be delivered within the kidney care setting. That makes it a liaison topic before a nephrology one, and no screening approach, treatment sequence or dosing rule appears here.

5 · Depression, demoralisation and adjustment in cancer: screening and treatment

Depression, demoralisation and adjustment in cancer is owned by the Consultation-liaison, geriatric and transcultural psychiatry notes. That chapter carries the distinction between demoralisation and depressive disorder, and the treatment sequence; this chapter carries neither, so a candidate revising how demoralisation differs from depression should not look here. A screening question or brief instrument for depression in advanced disease is stated by neither chapter, and that is a gap in the series rather than a section you have failed to find, so supply none from memory.

6 · Palliative care psychiatry: delirium at end of life, desire for hastened death and existential distress

Palliative care psychiatry is owned by the Consultation-liaison, geriatric and transcultural psychiatry notes, with delirium as a syndrome owned by the Stroke, TBI, delirium and focal brain syndromes notes and the doctrine of capacity owned by the Mental health legislation and forensic psychiatry notes. This chapter says only that capacity is decision-specific and time-specific, and states no approach to sedation, existential distress or a request for hastened death. What changes about delirium once treatment intent has changed, including how far to investigate and how antipsychotic choice differs in the final phase, is stated by neither owner and is a gap in the series rather than a section you have failed to find.

7 · Steroid-induced psychiatric disturbance: mania, depression, psychosis and management

Steroid-induced psychiatric disturbance is owned by the Endocrine and metabolic psychiatry notes. It is the row that follows a drug rather than a setting, so its dose relationship, its onset window and its management sequence all sit in that chapter and none is stated here. The chemotherapy and interferon section of this chapter carries no steroid limb either, by owner ruling, so there is exactly one place to look.

TRAP **Revising from the routing line as though it were an abstract.** A cross-reference is an address, written to be too short to revise from. The failure mode is a candidate who says depression after infarction matters to prognosis but cannot say what it does to it, or who knows corticosteroids cause psychiatric disturbance but cannot say when. Both collapse on the follow-up.

IN INDIA

No India-specific row is held for this section, and none has been invented. The Indian point is one of study order rather than epidemiology. Each of these constructs is by definition a liaison presentation: the patient is admitted under another team, and the psychiatrist arrives on request, often late and with a referral question written by somebody unsure what to ask. The chapters that own them also carry the referral, capacity and communication material, so read them before a liaison posting starts.

GOOD TO REMEMBER

At the bedside the question is not what you remember but which chapter answers the referral in front of you. Cardiac ward, intensive care, dialysis unit, oncology ward and palliative care unit all route to the liaison chapter. A psychiatric disturbance that began after a corticosteroid was started routes to the endocrine and metabolic chapter. Knowing that split in advance beats a half-remembered paragraph about any of them.

These constructs are named here, routed here and taught nowhere in this chapter, so that each is taught once, in full, where it is owned.

8 · Psychiatric aspects of heart failure, arrhythmia and cardiac devices

The failing heart produces most of the somatic symptoms a depression scale counts, so cardiac psychiatry is decided by history and by stage of illness rather than by a symptom total. Drug treatment is deliberately absent: prognosis and antidepressant choice in cardiac disease are taught with depression and anxiety after myocardial infarction and in coronary artery disease, and organ-adjusted dosing with the psychopharmacology notes. What is examinable here precedes any prescription.

8.1 The overlap that defeats a symptom count

Advanced heart failure and a depressive episode share almost the entire somatic side of the syndrome. Breathlessness on exertion, fatigue, disturbed sleep, poor appetite, weight change, psychomotor slowing and loss of sexual interest each follow from a failing circulation, from the treatment burden it imposes, or from a mood disorder, and none of them separates the possibilities alone. A scale that counts them inflates where accuracy matters most; a clinician who discounts them all misses the depression that is present.

Weight the **cognitive and affective items** and interrogate the somatic ones instead of scoring them. Anhedonia in activities that make no physical demand is the most useful single question on a cardiology ward: a person too breathless to walk to the gate can still say whether the radio or a visitor gives any pleasure. Guilt, worthlessness, hopelessness out of proportion to the prognosis actually given, and the thought that others would be better off without them are produced by mood.

Timing is the strongest discriminator on history. If breathlessness and oedema improved with cardiac treatment while mood, interest and outlook stayed flat, the mood state has separated from the haemodynamic state and is behaving as an illness in its own right. If low mood arrived with a decompensation and lifted as it settled, an adjustment reaction reads better. Nocturnal waking divides the same way: waking breathless and needing to sit upright points at the circulation, while waking in the early hours, alert, with the day at its worst on rising, points at mood.

In adults with confirmed heart failure, **people with heart failure and depressive symptoms had poor health-related quality of life**, so the depressive syndrome tracks with how the person actually lives rather than with a score alone. That finding carries a direction and a population and no magnitude, so its use is to justify looking, and to justify treating what is found.

RULE **Attribute by timing and by course, never by the symptom itself.** Any single somatic symptom in heart failure is uninformative. What discriminates is whether it moved with the cardiac state or independently of it, and only a sequential history answers that. Put the cardiac events, the treatment changes and the mood changes on one timeline and the attribution usually declares itself.

8.2 Stage of illness as the organising question

Heart failure is not one psychiatric situation, and the commonest error in a viva is to describe it as though it were. The task, the assessment priority and the realistic intervention all change with the stage.

STAGE	THE PSYCHOLOGICAL TASK	WHAT THE ASSESSMENT MUST REACH
New diagnosis	Absorbing a prognosis and a changed identity, often while still physically well	What they were told, what they believe it means, and who else heard it
Stable on treatment	A daily regimen, fluid and salt restriction, weighing, symptom diaries	Adherence, the cost of the regimen, and whether the effort has become hopeless rather than tiresome
Recurrent decompensation	Repeated admission, unpredictability, lost confidence in the body	Demoralisation, fear of dying in an episode, sleep, the carer's own state
Advanced illness	Anticipated dying, dependence, the meaning of continued treatment	Depression as distinct from a considered response to a real prognosis

The hand-offs are explicit. Demoralisation as a construct, and the screening and treatment of depression and adjustment in a life-limiting diagnosis, are taught in the organ-interface routing section, which also owns palliative care psychiatry and the desire for hastened death. Delirium in the decompensated or dying cardiac patient belongs to the stroke, traumatic brain injury and focal brain syndromes notes, including the bedside instruments.

Cognition is examined rarely and matters constantly. Where attention and processing speed are slowed in advanced disease, a crowded regimen and a consent conversation are at risk, and the risk peaks during a decompensation. Since **heart failure with depressive symptoms ran with poor health-related quality of life**, the stage review is a functional review as much as a diagnostic one.

8.3 Arrhythmia in the psychiatric clinic, and the separation from panic

Arrhythmia reaches psychiatry along several routes: a paroxysmal rhythm disturbance already called panic, an anxiety syndrome that grew out of a documented arrhythmia, a rhythm disturbance and an anxiety disorder present together, and the causes shared by each, which include alcohol, stimulants, thyroid disease and sleep deprivation. The separation is made on history.

Onset and offset carry most of the information. A paroxysmal rhythm disturbance tends to start and stop abruptly, sometimes with a trigger such as a change of posture, and the racing heart precedes the fear. A panic attack builds to a peak over minutes and settles more gradually, fear and bodily symptoms rise together, and catastrophic thinking about the heart is present at the peak. Episodes that wake a person from sleep, that carry syncope, or that occur on exertion rather than at rest shift the balance towards a cardiac cause.

TRAP **Diagnosing panic disorder in a patient whose rhythm has never been recorded during an episode.** They overlap almost completely at the level of reported symptoms, and the label is sticky: once panic is written, every later episode is read through it and the rhythm is never captured. The requirement is an ambulatory or event recording during a typical attack, and the formulation should state whether one exists.

The reverse error is commoner in cardiology. A person with a documented rhythm disturbance and an anxiety disorder is often told the palpitations are anxiety, which is heard as an accusation and produces exactly the hypervigilant self-monitoring that amplifies the next episode. The workable formulation is that the rhythm is real, the fear is real, and each raises the salience of the other.

8.4 Rhythm procedures and the peri-procedural window

Atrial fibrillation is an episodic, bodily felt and unpredictable illness, and unpredictability is the property that generates anticipatory anxiety in any organ system. The load in the group referred for rhythm procedures is substantial: **depressive and anxiety disorders are highly prevalent among patients with atrial fibrillation (AF) undergoing pulmonary vein isolation (PVI)**.

Because **depressive and anxiety disorders are highly prevalent in the group referred for pulmonary vein isolation**, screening belongs before the procedure rather than after a disappointing result. A person anxious before ablation who remains symptomatic after it is easily written off as a treatment failure when the rhythm record is clean. Naming it in advance changes how a residual symptom is read: as an expected part of recovery rather than as proof the procedure failed.

Symptom-rhythm concordance is the idea underneath this. Do the symptoms the patient reports coincide with the rhythm the record shows? Concordance supports a rhythm intervention; discordance says something other than the rhythm is generating the experience, usually interoceptive vigilance, which responds to psychological work rather than to further ablation.

8.5 Living with an implanted defibrillator

An implanted defibrillator is unlike any other treatment a psychiatrist meets, because its therapeutic action is an aversive event delivered without warning from inside the body. The psychopathology follows from that: the therapy hurts, its timing is unpredictable, and it cannot be escaped by leaving the situation. A sudden painful event paired with whatever the person happened to be doing is all ordinary conditioning needs.

Shock-related anxiety is anticipatory: continuous apprehension about the next discharge, with hypervigilant monitoring of the pulse, the chest and any sensation that preceded a previous one. **Avoidance** is the disabling part and the part most often missed, because the person rarely volunteers it. Exertion, sexual activity and driving are avoided in case they provoke the device, being alone in case a discharge occurs with nobody present, and public places because a discharge witnessed by strangers humiliates as well as frightens. After repeated discharges in a short period, intrusive recollection, hyperarousal and phobic avoidance can take a post-traumatic form. Some

patients describe a **phantom shock**, the vivid experience of a discharge the device record shows did not happen.

The household is part of the presentation. Partners commonly witness the discharge, and witnessing what looks like a resuscitation produces its own distress, its own hypervigilance and its own protective restriction of the patient's activity, so a couple can arrive at a shared avoidance that neither of them has named. Hence the recommendation that **psychological assessment, monitoring, and therapy should be offered to ICD patients and their partners as part of routine care**. Each noun does work: assessment at implantation and after any discharge, monitoring because a person well at the device clinic one month may be housebound the next, therapy because detection alone is not a service, and partners because the couple is the unit of care.

RULE **Ask what the person has stopped doing since the device was implanted.** Anxiety is reported readily and avoidance is volunteered almost never, yet avoidance is what removes exercise capacity, work, driving, intimacy and social life. Put the question to the patient and to the partner separately.

TRAP **Reading shock anxiety as simple reassurance-seeking and answering it with reassurance.** Reassurance relieves for minutes and teaches the person that the fear requires an external answer, so the checking increases. The device is genuinely capable of firing, so honest information plus a rehearsed plan for what to do when it fires converts an unbounded threat into an event with a script.

8.6 The device at the end of life, and the deactivation conversation

As cardiac disease advances, the shock function can fire during the dying process, with no prospect of restoring a life the person is already leaving. Deactivating the shock therapy is therefore a planned conversation, held while the person is well enough to hold it and recorded where the treating teams can find it. Psychiatry is asked in for three reasons: capacity, whether depression is colouring the decision, and how to hold the discussion with a family who may hear deactivation as abandonment.

Capacity here is decision-specific and time-specific, so the question is whether this person can make this decision now, and the answer can change with the delirium, the hypoxia or the sedation of an acute episode. The doctrine and the statutory test belong to the mental health legislation notes. What psychiatry adds is the mood assessment underneath: a wish to stop a burdensome therapy late in an illness can be a considered choice, whereas the same wish carried on pervasive hopelessness and a conviction of being a burden requires that the depression be treated and the question asked again.

Families frequently hear deactivation as switching the device off altogether, or as an act that causes the death. What deactivation covers is for the cardiology and device team to answer explicitly, in the family's own language; the psychiatric contribution is to ensure the question is asked and the answer understood.

8.7 Non-drug management, and who does what

Everything here is deliverable without a prescription, which matters where the cardiac regimen is already crowded.

- A written plan for what to do if a discharge occurs or an episode starts: rest, position, who to telephone, when to attend, and when the device record will be checked.
- A graded return to activity, agreed with the cardiologist in writing and shared with the partner. Avoidance is reversed by activity that has been sanctioned; encouragement without sanction moves the conflict into the household.
- Cognitive-behavioural work aimed at the maintaining cycle: catastrophic interpretation of a bodily sensation, checking, avoidance, and the relief that reinforces it. Interoceptive exposure needs explicit cardiology agreement first.
- Paced breathing and applied relaxation, taught as skills for the episode rather than offered as a cure.
- Couple-inclusive sessions as standard rather than as an extra, since **psychological assessment, monitoring and therapy are recommended for defibrillator patients and their partners as part of routine care**.
- Coordination with cardiac rehabilitation where it exists, and with the device clinic in every case, including a question about remote monitoring: it reassures some patients and feeds the vigilance of others.

Sleep is where the cardiac and psychiatric arms of the plan meet. Orthopnoea, nocturia from the regimen, fear of dying in sleep and anticipatory fear of a night-time discharge each fragment the night and each has a different remedy, and asking what wakes the person, and what they do next, separates them faster than any sleep questionnaire.

What belongs elsewhere, named: antidepressant choice, prognostic effect and cardiac safety sit with depression and anxiety after myocardial infarction and in coronary artery disease; organ-adjusted dosing with the psychopharmacology notes; the acutely stress-precipitated cardiomyopathy that mimics infarction with the stress-heart interface section; delirium with the stroke, traumatic brain injury and focal brain syndromes notes; palliative care psychiatry with the organ-interface routing section.

Timing, not the symptom WHAT DECIDES ATTRIBUTION	**Stage, not the label** WHAT ORGANISES THE ASSESSMENT	**Recorded, not assumed** WHAT THE PANIC LABEL NEEDS FIRST	**Avoidance, and the partner** WHAT THE DEVICE INTERVIEW MUST REACH

Each tile carries a decision rather than a fact, and not one of them can be settled from the symptom in front of you: attribution comes from the timeline, priority from the stage, the panic label from a recording, and the device interview from the partner as well as the patient.

IN INDIA

Device therapy is largely an out-of-pocket purchase, the device clinic may be a day's travel away, and the family rather than the individual usually carries the decision and the follow-up. The couple-inclusive element of the recommendation is therefore easier to deliver here than in services built around the individual patient, while longitudinal monitoring is harder, because the person is seen when they can travel rather than when they are due. A written plan the family can act on at home does more work here than a scheduled review. What this section does not hold is worth stating: no Indian prevalence figure for depression in heart failure, no Indian count of implanted devices, no Indian cost figure for defibrillator therapy.

GOOD TO REMEMBER

A short set of questions does most of the ward work: what has changed since the last cardiac event and in what order, what the person has stopped doing since then, and who else was in the room when the prognosis or the device was explained. Attribution, avoidance, and the partner who is carrying the other part of the distress.

The plate lays the limbs of the assessment side by side: the shared somatic symptoms with the questions that separate them, the arrhythmia presentations and where each routes, the psychopathology of living with a device, and the non-drug management that follows. Prescribing is absent from it by design.

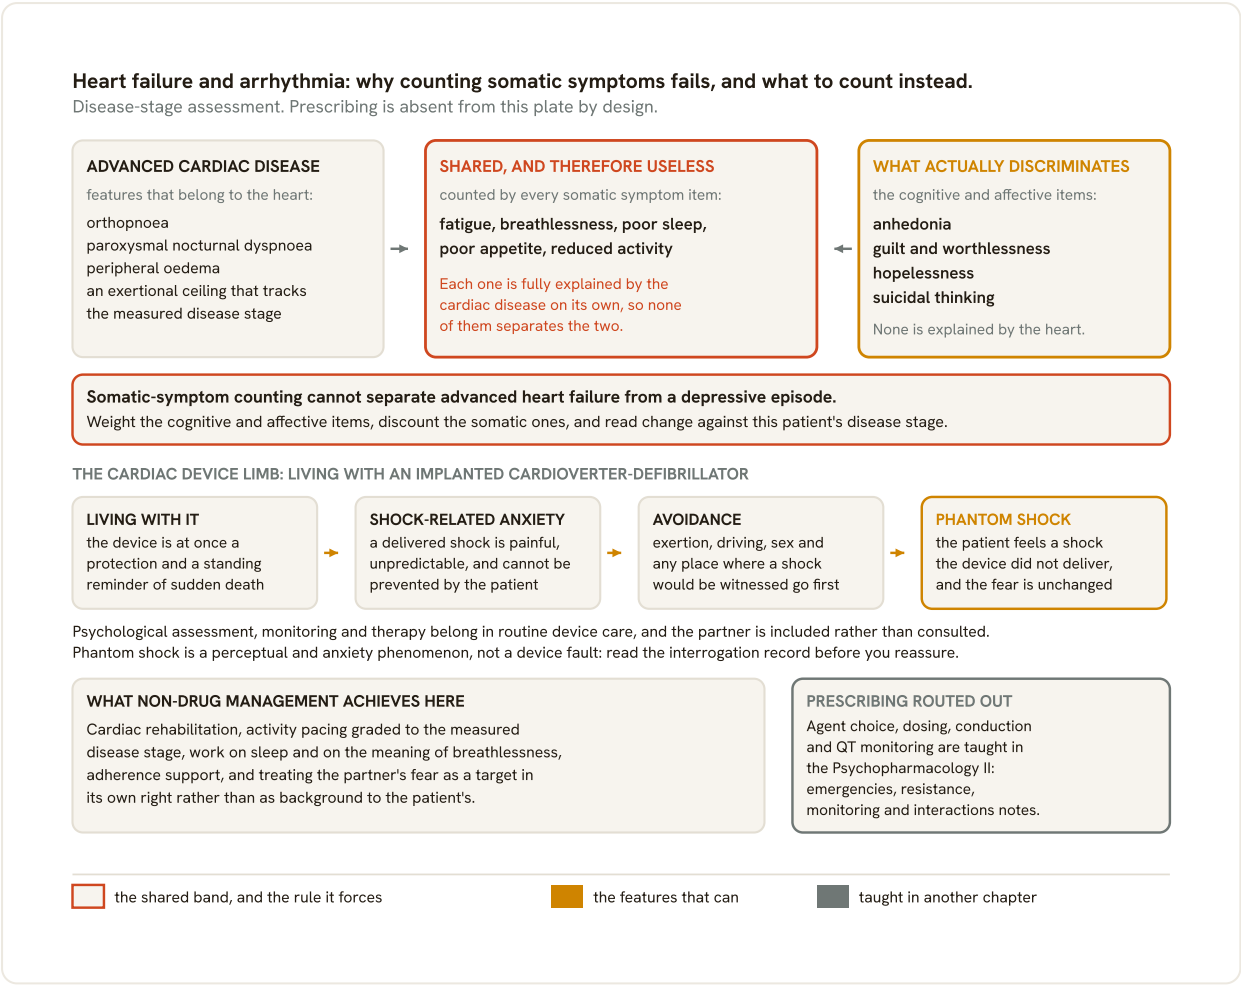


Fig 40.2 Disease-stage assessment in heart failure and arrhythmia, built around the symptom-overlap problem. The centre panel holds the five features a somatic symptom count would score, and every one of them is fully produced by advanced cardiac disease on its own, which is why an instrument weighted toward somatic items inflates in this population and a candidate who counts them reaches the wrong answer confidently. The discrimination lives in the right-hand panel, in items the heart does not generate. The device limb runs on its own logic: the implanted cardioverter-defibrillator protects and threatens at the same time, a delivered shock is painful and cannot be prevented by the patient, avoidance follows, and phantom shock closes the loop by producing the fear without the event. Prescribing is deliberately absent. From memory you should be able to name the five shared features, the four discriminating ones, and the four steps of the device limb, and say why the first list cannot settle the diagnosis.

> Attribute by timing and course, ask what the person has stopped doing since the device was implanted, and treat the couple, since psychological assessment, monitoring and therapy should be offered to defibrillator patients and their partners as part of routine care.

9 · Psychiatric aspects of COPD, asthma and chronic respiratory failure

Breathlessness and panic recruit the same body, and the patient living through either one cannot reliably tell you which of them started. Respiratory psychiatry is examined on that ambiguity and on what a clinician does about it at the bedside. This section installs four things: a working discrimination between the dyspnoea of obstructive airways disease and the dyspnoea of a panic attack; the two specialty-guideline positions this chapter can anchor to a respiratory source; the behavioural levers that move outcome, which are referral to rehabilitation and inhaler adherence; and an honest declaration about chronic respiratory failure, where a specialty

statement was sought and not obtained. Cardiac material is not taught here. Heart failure, arrhythmia and cardiac devices have their own section, and the psychiatric sequelae of critical illness, including post-intensive-care syndrome, sit in the merged organ-interface routing section.

9.1 The discrimination this section exists for

A patient who is breathless, frightened, tachycardic, tremulous and unable to complete a sentence has produced a presentation that two very different processes can generate. Airflow limitation produces breathlessness, and breathlessness produces fear; fear produces hyperventilation, and hyperventilation produces breathlessness. The loop closes in both directions, which is why the question is rarely which one the patient has, and almost always which one is driving the episode in front of you.

Several features shift the probability without settling it. The temporal shape is the most useful single discriminator: a panic attack crescendos and then subsides over a bounded period, whereas an exacerbation of airways disease climbs and stays. The relation to exertion matters, because bronchoconstriction and deconditioning both load onto effort while a panic attack characteristically arrives at rest and out of context. Objective airflow limitation, audible wheeze, a fall in oxygen saturation and a response to a reliever inhaler point one way; circumoral and peripheral paraesthesiae, carpopedal spasm and the light-headedness of hypocapnia point the other, since an obstructed patient usually cannot blow off enough carbon dioxide to produce them. The content of the thought is informative and is often the most accessible part of the history: the catastrophic cognition in panic is about dying or losing control now, and it recurs in the same form.

None of this licenses a clean split, and the examinable position is that it does not need to. A patient with obstructive lung disease who develops recurrent panic attacks has both conditions, each worsening the other, and treating only one leaves the disability where it was. **Anxiety sensitivity**, the tendency to read bodily sensations as dangerous, is the variable that couples them, and it is the one a psychiatrist can treat.

RULE **The dangerous direction of error is treating airflow limitation as anxiety, not the reverse.** Calling a panic attack an exacerbation costs an unnecessary course of treatment. Calling an exacerbation a panic attack can cost a life. Where the two cannot be separated in the room, treat the airway first, interrogate the psychiatry afterwards, and record why.

9.2 Why mood and anxiety cluster in obstructive lung disease, and what they cost

The association is not a coincidence of two common conditions sharing a population, and a candidate who says only that will be pressed. Several distinct pathways converge, and naming them separately earns the marks.

The behavioural pathway is the strongest and the most modifiable. Breathlessness on exertion is aversive, so the patient reduces exertion; reduced exertion produces deconditioning, and deconditioning lowers the threshold at which the next effort produces breathlessness. The spiral narrows the patient's world by degrees and delivers the ingredients of depression, which are loss

of activity, loss of role and loss of contact. From the anxiety side it is avoidance learning: exertion is paired with a frightening sensation, and the patient stops.

Beyond that sit contributions this chapter can name but does not quantify. Tobacco use is a shared risk factor for the lung disease and for several psychiatric conditions, and it means a proportion of these patients carry a nicotine dependence any plan must address. Chronic hypoxaemia and hypercapnia affect attention, processing speed and daytime alertness, and disrupted nocturnal breathing fragments sleep, which independently degrades mood. Systemic corticosteroid exposure is common here and carries its own psychiatric consequences; that construct is owned elsewhere and is named only as a hand-off. Finally, the disease is progressive, visible and frightening with a foreseeable trajectory, and demoralisation in that setting is a reasonable response as well as a clinical problem.

The cost is not only distress. Comorbid mood and anxiety symptoms track with poorer adherence, more frequent presentation to acute services and worse functional status at any given level of measured airflow limitation. No prevalence figure, effect estimate or readmission proportion is held by any claim row available here, so none is stated: carry the direction, and say openly that the magnitude was not obtained.

TRAP **Deciding that low mood in advanced lung disease is understandable, and therefore leaving it alone.** Comprehensibility is not a diagnostic exclusion, and this reasoning would not be accepted for any other comorbidity. It matters more than usual here because the depressive component is what removes the patient from the activity, the rehabilitation and the adherence that would change the trajectory. An understandable depression and a treatable depression are the same depression.

9.3 Pulmonary rehabilitation, and the referral statement to memorise

Pulmonary rehabilitation is a structured programme of supervised exercise training combined with education and self-management support for patients with chronic respiratory disease. It belongs in a psychiatry chapter because it acts on the behavioural spiral above: it is the intervention most likely to reverse the deconditioning, avoidance and social withdrawal carrying much of the psychiatric burden here.

The single most examinable statement this section owns concerns who gets referred. The specialty guideline position is that **coexistent symptoms of anxiety and/or depression in patients with COPD should not preclude referral to pulmonary rehabilitation**. That sentence is worth learning in its shape as well as its content, because it is written as a prohibition on exclusion rather than as an encouragement to refer.

It is written that way because the pressure in real services runs the other way. A patient who appears low, unmotivated or unlikely to complete a course is exactly the patient a stretched programme is tempted to screen out, and the screening is usually informal, done by the referrer, on a judgement about probable engagement rather than by any stated criterion. The guideline blocks that move, and the logic is worth stating in a viva: the psychiatric symptoms are in

substantial part a product of the very deactivation the programme treats, so using them to withhold it removes the intervention from the group with most to gain.

Two boundaries keep the statement honest. It concerns referral and not completion, so supporting attendance, which is where anxious and depressed patients genuinely fall away, remains a clinical task. And no effect size, completion rate or symptom change accompanies the recommendation in the row this chapter holds, so no magnitude is claimed. What a candidate can say without qualification is that coexistent anxiety and depressive symptoms are not a reason to withhold referral to pulmonary rehabilitation, and that the psychiatrist's contribution is often to say so to the referring team.

RULE **Psychiatric comorbidity is an indication to support the referral, not a filter applied before it.** When a respiratory team asks whether a low or anxious patient is suitable for rehabilitation, the guideline-anchored answer is that coexistent symptoms of anxiety and/or depression should not preclude referral. The useful follow-up question is not whether to refer, but what must be arranged so the patient reaches the second session.

9.4 Adherence and technique, which is the behavioural half of disease control

Inhaled treatment is the one therapeutic area where a psychiatrist's ordinary skills change a respiratory outcome, because almost every failure of inhaled treatment is behavioural rather than pharmacological. Two distinctions organise the assessment. The first is between **unintentional non-adherence**, where the patient intends to take the treatment and does not manage it, and **intentional non-adherence**, where the patient has decided against it. The second is between not taking the drug and not delivering it, since a patient with poor inhaler technique is fully adherent and receiving very little.

Psychiatric symptoms map onto these categories specifically rather than generally. Depressive symptoms produce forgetting, loss of routine and the hopeless reasoning that treatment will not help anyway. Anxiety produces overuse of the reliever inhaler and a dependence on it that masks deteriorating control and intermittently reinforces using it again. Beliefs about harm, particularly the fear that inhaled steroid preparations are dangerous or habit-forming, produce quiet intentional non-adherence the patient will not volunteer. Cognitive impairment and manual difficulty produce technique failure independent of all of this.

FAILURE MODE	WHAT IT LOOKS LIKE IN CLINIC	WHAT ACTUALLY CHANGES IT
Technique failure	Correct prescription, exhausted canister, poor control, patient reports full adherence	Watching the patient use the device, then re-teaching or changing the device
Forgetting and lost routine	Erratic use, worse in the morning, tracks with mood and with disorganised days	Anchoring the dose to an existing fixed daily act, plus treating the mood state
Reliever overuse	Frequent reliever use, escalating anxiety, preventer untouched	Naming the cycle explicitly, then treating the anxiety alongside the airway
Belief-driven refusal	Articulate reasons, often about steroids, offered only when asked directly	Eliciting the belief without correcting it first, then addressing the specific fear

The instruction that follows is single and easy to examine: ask the patient to demonstrate, do not ask the patient to report. A demonstration takes less time than the conversation about escalating treatment that follows from assuming the technique was adequate.

Dyspnoea against panic: six discriminators, and the band where they genuinely overlap.
Neither outer column settles the question alone. The middle column is the reason it stays hard.

FEATURE	DYSPNOEA OF COPD OR ASTHMA	WHERE THEY OVERLAP	PANIC ATTACK
ONSET	over minutes to hours, from a baseline that was already limited	both begin with breathlessness and fear, in the same body	abrupt, peaking within minutes from an unremarkable baseline
TRIGGER	exertion, cold air, allergen, infection or a missed preventer	exertion can start either, and fear can follow either	often at rest and out of context, or cued by fear of an attack
TIME COURSE	climbs, and stays up until the airway is treated	a long attack and a short exacerbation look alike	crescendos, then subsides over a bounded period
WHAT RELIEVES IT	a reliever inhaler, and treating the airway itself	reliever overuse eases both for a while, and masks both	slowed breathing and time: the attack is self-limiting
WHAT THE PATIENT DOES	drops exertion, then drops the activities that needed it	avoidance narrows the life the same way from either cause	escapes the situation, carries the inhaler as a safety object
WHAT THE OBSERVER SEES	wheeze, accessory muscle use, falling oxygen saturation	tachypnoea, tachycardia, tremor and visible fear in both	circumoral and peripheral paraesthesiae, carpopedal spasm

The dangerous direction of error is treating airflow limitation as anxiety, never the reverse.
Where the two cannot be separated in the room, treat the airway first, interrogate the psychiatry after, and record why.

LEVER ONE: PULMONARY REHABILITATION
Coexistent symptoms of anxiety and/or depression in patients with COPD should not preclude referral to pulmonary rehabilitation. It is written as a prohibition on exclusion, not an encouragement.

LEVER TWO: INHALER ADHERENCE AND TECHNIQUE
Separate not taking the drug from not delivering it, and unintentional from intentional non-adherence. Then ask the patient to demonstrate the device rather than to report on it.

ROUTED OUT Steroid-induced psychiatric disturbance is taught in the Endocrine and metabolic psychiatry notes.
ROUTED OUT Agent choice, dosing and monitoring are taught in the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes. This plate names no agent and prints no dose.

Fig 40.3 Separating the dyspnoea of obstructive airways disease from a panic attack, with the overlap drawn rather than glossed. The two outer columns are the features that shift probability, read across six axes rather than argued one at a time. The coral column is the honest part: breathlessness and fear recruit the same body, the loop closes in both directions, and a patient with obstructive lung disease who also has panic attacks has both conditions, each worsening the other. That is why the examinable question is almost never which one the patient has but which is driving this episode and in what proportion. The two levers below are where a psychiatrist changes a respiratory outcome, and both are behavioural rather than pharmacological. From memory you should be able to run the six axes in either direction, state the asymmetry of risk in the rule box, and give the referral sentence and the demonstration instruction verbatim.

9.5 Asthma in adolescence, and the association that has to be stated carefully

Asthma carries a second and separate examinable position, and it is age-specific. The specialty guideline records that **asthma in adolescence is associated with an increased likelihood of major depression, panic attacks and anxiety disorder**. Three things about that sentence are worth holding. It names adolescence rather than asthma at any age. It names three outcomes rather than a general elevation in distress. And it states an association, so a candidate who converts it into a causal claim in either direction has overstated the source.

The direction genuinely is uncertain and saying so is the stronger answer. A chronic, visible, sometimes frightening condition that constrains sport, sleep and social life at the stage where all three matter most is a plausible cause of low mood. Anxiety, equally, worsens the perception of breathlessness and the behaviours that control it. Shared factors sit underneath both, including disrupted sleep, corticosteroid exposure and family stress.

The consequence for practice is concrete. Because adolescent asthma is associated with an increased likelihood of major depression, panic attacks and anxiety disorder, mood and panic symptoms are worth asking about routinely in this group rather than waiting for them to be raised, and mood and anxiety are candidate contributors to poor asthma control rather than a separate matter for later. The overlap with panic is awkward because both under-perception and over-perception of airflow limitation occur: the adolescent reporting severe symptoms may be obstructed, panicking, or both, and the one reporting little may be seriously obstructed.

TRAP **Labelling an adolescent's attack as panic and thereby delaying respiratory treatment.** The association above makes panic more likely in this group, which is exactly why the label becomes available and exactly why it is dangerous. An adolescent with asthma presenting with an acute frightening episode is treated for the airway on clinical grounds, and the psychiatric formulation is completed once the episode has settled. Offering reassurance as the first move in an acute presentation answers the question badly, however well the comorbidity is discussed afterwards.

9.6 Chronic respiratory failure, and the gap this chapter declares

The third condition in the section's title is the one for which no guideline statement was obtained, and the honest position has to be stated rather than papered over. A specialty-guideline statement specifically linking chronic respiratory failure with psychiatric sequelae, with psychiatric assessment, or with psychiatric management was sought for this chapter from the respiratory bodies whose statements would carry such a recommendation, being the national thoracic guideline body and the international obstructive lung disease initiative. It was not obtained in a form this chapter can bind. That absence describes what this chapter opened. It is not evidence that no such statement exists, and it must not be converted into a claim that the specialty is silent.

What a candidate does without it is reason from general principles and say that this is what they are doing. Chronic hypercapnia produces daytime somnolence, headache, impaired attention and, as it advances, confusion, any of which can present first as an apparent psychiatric change. Nocturnal hypoventilation fragments sleep, which degrades mood independently. Long-term oxygen therapy and domiciliary ventilation impose visible dependence on a machine and reorganise family life around the equipment. The loss of control is the psychological core, and interventions that return any control to the patient tend to help.

Two boundaries apply. Acute confusional states here are assessed and managed under the delirium material in the stroke, traumatic brain injury and focal brain syndromes notes, and no bedside delirium instrument, rating scale or antipsychotic choice is taught in this section. And a patient declining domiciliary ventilation raises a capacity question; capacity is decision-specific and time-specific, and everything beyond that, including the statutory test, belongs to the mental health legislation notes.

IN INDIA

No figure for the availability, distribution or coverage of pulmonary rehabilitation programmes in India is held by any claim row available to this chapter, so none is offered. What can be said is what the referral position implies for a service that has a programme and for one that does not. Where a programme exists, coexistent anxiety and depressive symptoms are not grounds to withhold referral to pulmonary rehabilitation, and a psychiatrist asked to assess a low patient before referral has a guideline-anchored answer for the respiratory team. Where no programme exists, which is the position across much of the country outside larger centres, the deliverable residue is still the active ingredient: a graded, written, supervised increase in activity that a physiotherapy service, a nursing team or an informed family member can sustain, with education about the breathlessness cycle. The gap in a district service is usually the structure and the supervision, not the exercise.

9.7 What this section hands off, and does not teach

Corticosteroid-induced psychiatric disturbance is common in this population and is owned elsewhere. It is named here as a differential whenever mood, sleep or mental state changes after a course of systemic steroid; the mania, depression, psychosis and management content sits with the corticosteroid owner in the endocrine and metabolic material and on the chapter's routing plate.

Prescribing in respiratory disease routes to the psychopharmacology notes without exception. Sedative burden and respiratory depressant risk are why a psychiatrist is cautious here, and the drug-by-drug reasoning, the comparative rankings and every dose belong there, so none is restated. The same applies to the effect of tobacco smoke on the metabolism of certain psychotropics: it is named as a hand-off, and neither its direction nor its magnitude is stated here.

Non-pharmacological treatment of the anxiety component follows ordinary principles with one adaptation. Interventions that deliberately provoke the feared bodily sensation require care when the airways are genuinely obstructed, because the sensation is not false. The workable version treats the interpretation and the avoidance rather than trying to show the sensation is harmless, and it runs alongside optimised inhaled treatment.

GOOD TO REMEMBER

On the ward, three moves cover most of the value a psychiatrist adds to a breathless patient. Watch the inhaler being used before believing anything about adherence. Ask what the patient thinks is happening during an attack, because the answer separates the catastrophic cognition of panic from the accurate fear of someone obstructed before. And ask whether the patient has been referred for rehabilitation, because anxiety and depressive symptoms are not a reason to preclude that referral and a psychiatric opinion is often what is used to justify not making it.

9.8 The answer, assembled

An examiner asking about psychiatric aspects of respiratory disease wants a structure rather than a list. The following sequence carries this section.

- State the discrimination first: breathlessness and panic share a physiology, the loop runs both ways, and the common answer in this population is both rather than either.
- Give the safe rule of error: treat the airway first when the two cannot be separated, because the costly mistake is calling obstruction anxiety.
- Explain the behavioural spiral of breathlessness, avoidance and deconditioning, since it accounts for much of the presentation and is what treatment targets.
- Deliver the referral position verbatim: coexistent symptoms of anxiety and/or depression in patients with COPD should not preclude referral to pulmonary rehabilitation.
- Deliver the asthma position with its age qualifier and its association wording: asthma in adolescence is associated with an increased likelihood of major depression, panic attacks and anxiety disorder, and say the direction was not established.
- Cover adherence and technique as a behavioural target, and demonstrate rather than ask.
- Declare the chronic respiratory failure gap openly: reason from hypercapnia, sleep, dependence and loss of control, and say that a specialty statement binding psychiatric management to chronic respiratory failure was sought and not obtained.

Both, not either	**Refer, not exclude**	**Watched, not reported**	**Declared, not inferred**
WHAT DYSPNOEA PLUS PANIC USUALLY MEANS	WHAT COEXISTENT LOW MOOD CHANGES	HOW INHALER ADHERENCE IS ASSESSED	WHAT TO SAY ON RESPIRATORY FAILURE

Each tile names a default that ordinary practice sets the other way, which makes the strip an audit of one's own clinic behaviour rather than a summary of the text.

PITFALL

Answering this item as a general essay on the psychological impact of chronic illness. The material is true and it earns very little, because it could have been written about any organ. What makes an answer respiratory is the dyspnoea and panic discrimination, the referral position on rehabilitation, the age-specific asthma association and the behavioural analysis of inhaled treatment.

Breathlessness and panic are usually both present; treat the airway first, refer for rehabilitation regardless of coexistent anxiety or depression, watch the inhaler being used, and ask every adolescent with asthma about mood and panic.

10 · Takotsubo cardiomyopathy and the stress-heart interface

Takotsubo syndrome is the plainest demonstration in medicine that an acute psychological event can produce a structural cardiac lesion, which is why it belongs in a psychiatry paper and not only in a cardiology one. An emotional or physical stressor is followed, usually within hours, by chest pain, breathlessness or collapse, by an electrocardiogram reading as an evolving anterior infarct, by rising biomarkers of myocyte injury, and by a left ventricle gone akinetic over

a territory larger than any single coronary artery supplies. The arteries are then found unobstructed, and the wall-motion abnormality resolves, which infarcted muscle never does. This section teaches the syndrome by mechanism rather than by table: the epidemiological specifics a resident would ordinarily quote are absent from this build's rows, and are declared rather than supplied from memory.

Several neighbouring topics are named here and taught elsewhere, and knowing the boundary is itself examinable. Depression and anxiety after myocardial infarction, and safe antidepressant choice in established coronary disease, belong to the organ-interface routing section. Heart failure, arrhythmia and implanted cardiac devices have their own section. Delirium on the coronary care unit sits with the notes that own delirium, and the psychological sequelae of critical illness with intensive-care psychiatry. Organ-adjusted dosing is a psychopharmacology question. What remains is the mechanism, the differential and the liaison contribution.

10.1 The shape of the syndrome, and why a psychiatrist is called at all

The syndrome is defined by a pattern rather than by any single test, and each element carries weight. There is an acute regional abnormality of left ventricular contraction, extending beyond the distribution of one coronary artery, which makes an occlusive explanation implausible on anatomical grounds alone. There is no culprit obstruction able to account for it. And there is recovery, so the ventricle akinetic on admission contracts normally later. Transience is not a happy outcome appended to the definition, it is part of the definition, and why the diagnosis is partly retrospective.

The name describes the shape the ventricle takes at end systole in the commonest form, a rounded akinetic apex above a narrow hypercontractile base, resembling a narrow-necked pot. That image encodes the paradox the mechanism has to explain: the segment furthest from the base fails while the base works harder than usual.

A psychiatrist is called into this picture for a small number of reasons, and knowing them in advance determines what a useful consultation note contains. The antecedent was psychological and somebody wants it characterised. A psychotropic drug, a substance or a withdrawal state is suspected. The patient is distressed or agitated on a cardiac unit. Or the team, having found clean arteries, is uncertain what to tell the patient and what to arrange on discharge. Each is answerable. Reciting a cardiology description back to a cardiologist is not.

10.2 The catecholamine-surge account, taken in the order the account itself runs

The dominant mechanistic account is a **catecholamine surge** acting on myocardium, and it is best learned as a chain, because each link is separately testable and separately arguable.

The chain begins centrally. A stressor is appraised as threatening, or the body is placed under acute physiological load, and the central autonomic network drives the sympatho-adrenal axis. The adrenal medulla releases adrenaline into the circulation; cardiac sympathetic terminals release noradrenaline into the myocardial interstitium. These are separate deliveries with different distributions, which matters later: circulating adrenaline reaches the whole ventricle, whereas neural noradrenaline concentrates wherever innervation is densest.

The next link is the effect of that load on the myocyte. High catecholamine exposure drives calcium entry and calcium overload, and calcium-overloaded myocytes contract in a disordered, sustained way. The histological correlate described in this setting is **contraction band necrosis**, a pattern distinct from the coagulative necrosis of coronary occlusion, accompanied by oxidative stress, mitochondrial dysfunction and falling cellular energy supply. This is a toxic injury, not an ischaemic one, and the distinction is why the muscle can recover.

A parallel link is vascular. Catecholamine excess is associated with epicardial coronary spasm and, more importantly here, with **microvascular dysfunction**: impaired dilatation of the small vessels that determine actual myocardial perfusion. Perfusion can therefore fall while the large arteries remain widely patent, which reconciles a genuinely injured myocardium with a normal angiogram.

The most interesting link, and the one that best explains why the apex, is receptor level. Adrenoceptor density is described as higher towards the apex while sympathetic innervation is sparser there, leaving the apex disproportionately exposed to circulating adrenaline. The subtype predominating apically can also switch its intracellular coupling: at ordinary concentrations it signals through the stimulatory G protein and increases contraction, and at the concentrations reached in a surge it couples instead to the inhibitory G protein, reducing contraction. On this account the apex is not merely overwhelmed, it is actively switched off, and the switch is protective, sparing the myocyte a lethal calcium load at the cost of regional akinesia.

The final link is **myocardial stunning**. Stunned myocardium is alive but not contracting: the cells persist, the contractile apparatus is temporarily disabled, and function returns when the insult is withdrawn. Stunning is why a ventricle that looks infarcted on admission can look normal later, and it is the single concept that makes the whole syndrome coherent.

RULE **Learn the chain, not the label.** Central appraisal, sympatho-adrenal outflow, adrenal and neural catecholamine delivery, calcium overload with contraction band necrosis, microvascular impairment, apical receptor coupling switch, regional stunning, then recovery. A candidate who can walk that chain can answer any question asked about this syndrome, including those on variants and differential diagnosis, because each attacks a specific link.

10.3 What the account explains well, and where it strains

Good teaching of a mechanism includes its limits, and this one has clear limits. The apical receptor argument explains the commonest morphology elegantly. It explains it too well, because the syndrome also occurs with the akinetic segment midventricular, basal with a hypercontractile apex, and focal. A gradient anchored at the apex predicts apical disease; it predicts basal disease poorly. The honest reading is that the receptor gradient is one contributor within a more general catecholamine-mediated regional vulnerability, and that regional variation reflects unresolved individual differences in innervation, receptor distribution and microvascular reactivity.

The account also under-determines who develops the syndrome. Acute catecholamine surges are common in human life and this presentation is not, so a susceptibility term is required: something about vascular reactivity, hormonal environment, autonomic set point or genetic

variation separating the person who stuns from the far larger number who do not. Naming that term, and saying it is unresolved, is a better answer than asserting a single cause.

This is where a conventional account would supply a demographic profile, since the demographic skew of this syndrome is the specific most often reproduced about it. That figure is exactly what this build cannot supply: the rows available here carry no distribution by sex, no age distribution, and no proportion of acute coronary presentations accounted for by the syndrome. What can be taught instead is what such a figure would require: a consecutive unselected series taken for angiography with suspected acute coronary syndrome, an explicit diagnostic definition applied prospectively, and a stated catchment and period, so that numerator and denominator each mean something. Any figure quoted without those is a number without a population.

TRAP **Treating the receptor gradient as the whole explanation.** It is an attractive and memorable argument, and it accounts for the classical morphology only. Candidates who commit to it are then unable to explain the basal and midventricular variants, and an examiner who asks about variants is usually asking precisely this question.

10.4 Why the presentation is mistaken for an acute coronary syndrome

It is mistaken for an acute coronary syndrome because at first contact it is genuinely indistinguishable from one, which is a property of the syndrome rather than a failure of the clinician. The symptom is the same: chest pain, breathlessness, syncope or shock. The electrocardiogram carries ischaemic change in a territory, commonly ST-segment elevation anteriorly, later T-wave inversion and prolongation of the corrected QT interval. The biomarkers rise, because myocytes really are injured. The ventricle really is akinetic. Every element that ordinarily signals an occluded artery is present, and the artery is not occluded.

What separates them is therefore investigation and time, not acumen at the bedside. Angiography shows no culprit lesion. Imaging shows a wall-motion abnormality crossing coronary territories. Follow-up imaging shows recovery. None of that is available in the first minutes, which is why the correct initial management is the acute coronary syndrome pathway, and why a psychiatrist confident at the bedside that this is stress-related is confident about something they cannot yet know.

The psychiatric traps here run in opposite directions and each is serious. The first is re-labelling: the angiogram is reported clean, the word psychological enters the notes, and the presentation is reinterpreted as anxiety, panic or somatisation. Unobstructed arteries are compatible with a severely injured ventricle, with pulmonary oedema, arrhythmia and haemodynamic collapse: a clean angiogram is not a normal heart. The second is the mirror image: a recognisable emotional antecedent leads the team to a stress explanation while an atherothrombotic occlusion is in progress. A psychological trigger has no power to exclude an occlusion.

TRAP **Reading a clean angiogram as a psychiatric diagnosis.** The angiogram excludes an obstruction; it excludes nothing else. The commonest way this syndrome is mismanaged on a psychiatric liaison list is that the cardiac finding is treated as ruled out and the patient is handed over as an anxiety presentation, at which point nobody is monitoring a ventricle that is still stunned.

10.5 The psychiatric antecedents, and what can honestly be said about them

Antecedents fall into recognisable classes, and the classification here is conceptual rather than quantitative. Emotional antecedents include bereavement, acute interpersonal conflict, fear, grave news, an assault on status or reputation, and acute panic. Physical antecedents include surgery, acute severe illness, sepsis, seizure, respiratory failure and metabolic disturbance, engaging the same sympatho-adrenal pathway with no psychological content. A proportion of cases have no identifiable antecedent of either kind, which matters, because demanding one before entertaining the diagnosis will miss those patients.

A subtler point is valence. The mechanism is driven by arousal and catecholamine load, and arousal is not the exclusive property of aversive emotion, so intensely positive events have also been described as antecedents. This follows from the chain above and is not an anomaly needing separate explanation. What proportion of antecedents is emotional against physical, and what proportion is aversive against pleasurable, are quantitative questions with no grounded answer here.

Comorbidity is a different question from causation, and here this section does hold a grounded statement. **Psychiatric disorders, depression and anxiety among them, are common in patients with this syndrome**, and that is worth stating precisely for what it is: a statement of co-occurrence in the affected population. It establishes that the psychiatrist has a population-level reason to be involved rather than an incidental one.

What it leaves open is the direction of the relationship, and an examiner rewards a candidate who says so. Depression and anxiety may be an antecedent vulnerability, plausibly through sustained autonomic and hypothalamic-pituitary-adrenal dysregulation that alters the response to a surge. They may be a consequence, since surviving a frightening cardiac event with an uncertain explanation generates fear and low mood. They may reflect a shared substrate, with autonomic dysregulation producing each. Co-occurrence is compatible with all of those, and the proportion affected, with any temporal ordering, is a specific this build does not carry.

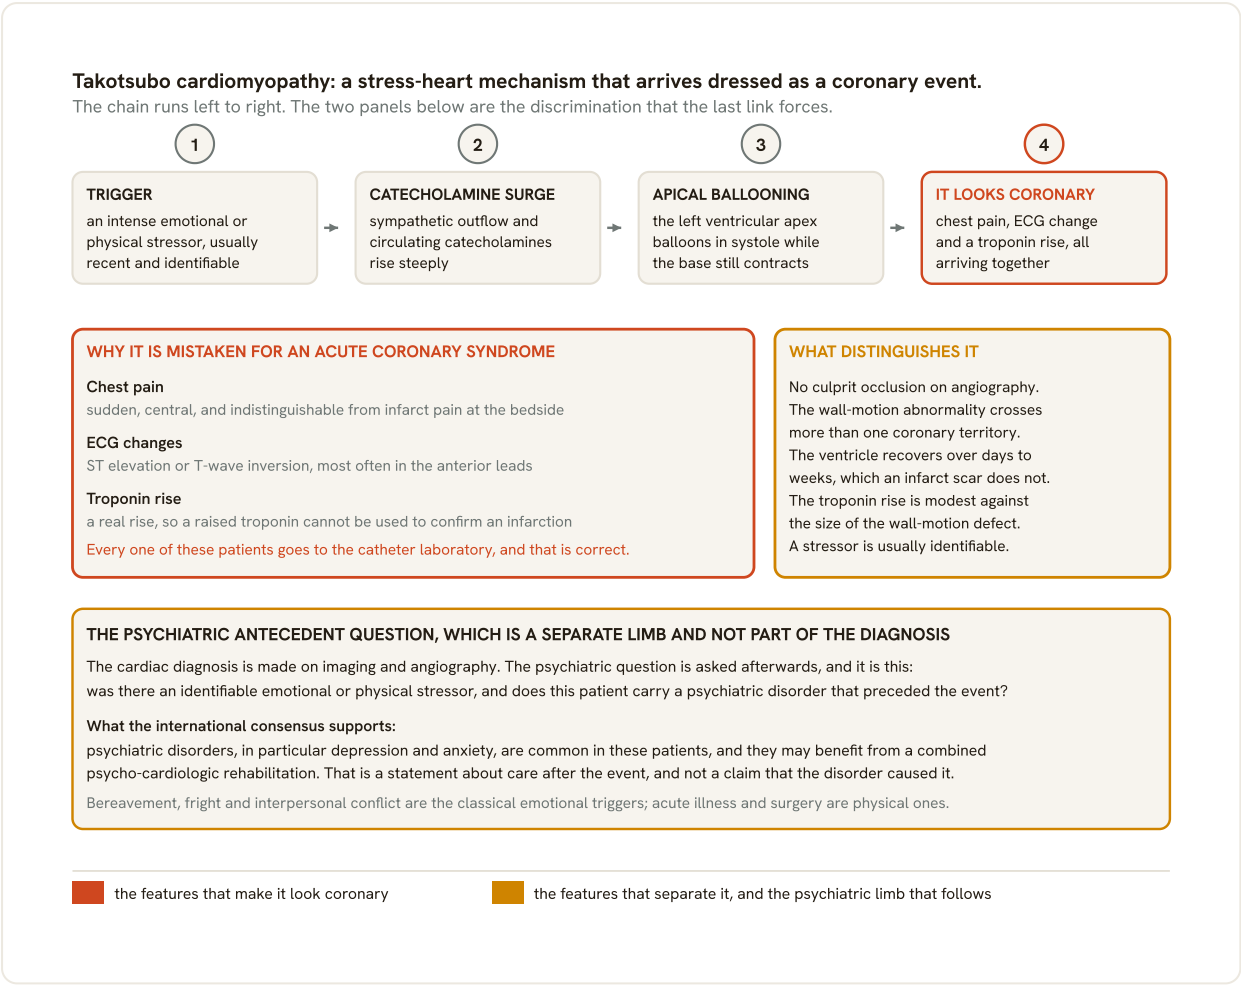


Fig 40.4 The stress-heart interface as a four-link chain, with the mimicry and the discrimination drawn side by side. The mechanism is the examinable spine: a trigger, a catecholamine surge, an apex that balloons in systole while the base still contracts, and a presentation that carries all three cardinal features of an acute coronary syndrome at once. That is why the mistake is not carelessness, and why the correct first move is still the catheter laboratory. What separates the two is the pattern rather than any single test: the wall-motion abnormality does not respect a coronary territory, no culprit occlusion is found, and the ventricle recovers in a way an infarct scar does not. The psychiatric limb is deliberately drawn apart from the diagnosis, because it is asked after the cardiac question is settled. From memory you should be able to give the four links in order, the three mimicking features, at least three discriminating ones, and the reason the antecedent question is a separate limb rather than a diagnostic criterion.

10.6 The differential of the surge, which is largely a psychiatric differential

Because the mechanism is catecholamine load, the differential of the antecedent is a list of states psychiatrists meet constantly, and this is where the liaison psychiatrist supplies information nobody else can obtain.

- **Exogenous catecholamine or sympathomimetic load.** Adrenaline given therapeutically, local anaesthetic containing adrenaline, and sympathomimetic substance use including stimulants and cocaine. The substance history establishes these.
- **Withdrawal states.** Alcohol and sedative-hypnotic withdrawal produce sustained autonomic arousal, and opioid withdrawal a lesser but real adrenergic load. A withdrawal syndrome is an endogenous catecholamine infusion the patient is unaware of receiving.
- **Iatrogenic surges in psychiatric treatment.** Electroconvulsive therapy produces a transient, deliberate and monitored autonomic surge, a recognised context in which this presentation

has been described. Management of that surge belongs to the treatment's own protocol.

- **Hyperadrenergic psychiatric emergencies.** Neuroleptic malignant syndrome, serotonin toxicity and catatonia with autonomic instability sustain adrenergic arousal for hours or days, a longer exposure than a discrete emotional trigger delivers.
- **Acute severe panic and acute stress reactions.** These are most likely to be dismissed as merely psychological on a medical ward, and the mechanism gives a reason to take them seriously as physiological events.
- **Organic causes with psychiatric presentations.** Phaeochromocytoma delivers precisely the surge the mechanism requires and may present with paroxysmal anxiety. Subarachnoid haemorrhage produces neurogenic stunned myocardium by an overlapping pathway, and prolonged seizure activity does the same. Each can reach a psychiatrist first.

Psychotropic drugs with cardiac effects, including corrected QT interval prolongation and the cardiovascular effects of stimulants, are a psychopharmacology topic, cross-referred rather than restated. Their relevance here is that the medication list is part of the surge history.

FIGURE CONVENTIONALLY QUOTED	WHAT WOULD HAVE TO BE MEASURED	WHAT THIS BUILD HOLDS
Frequency among suspected acute coronary presentations	Consecutive unselected patients angiographed for suspected acute coronary syndrome, prospective definition, stated catchment and period	Not grounded here
Sex and age distribution	Composition of a consecutive diagnosed series, read against the composition of the source population	Not grounded here
Emotional against physical antecedents	Antecedent classified prospectively by stated criteria in every case, including cases with none identified	Not grounded here
Time to recovery of wall motion	Protocol-defined repeat imaging at fixed intervals, with a stated definition of normalisation	Not grounded here
In-hospital complications and mortality	Pre-specified complication set, admission to discharge, on a consecutive series	Not grounded here
Recurrence	Cumulative recurrence over a stated follow-up period, with defined ascertainment	Not grounded here
Depression and anxiety in this population	Case ascertainment by a stated diagnostic method, at a stated point after the acute event	**Qualitative statement of commonness only**

10.7 What the psychiatrist on the liaison call actually contributes

The consultation has a small number of jobs, and doing them well is the difference between a note that changes management and a note that restates the referral.

The first is to characterise the antecedent properly. A cardiology record says stress. A psychiatric record says what happened, at what hour, with whom, what it meant to this person, and how long the arousal lasted before the chest pain began. That timed narrative is the only evidence anyone will have about the trigger, it can be obtained on the day and rarely afterwards, and no other member of the team is trained to take it.

The second is to establish the exogenous and withdrawal load, using the differential above. This is a diagnostic contribution, because a withdrawal syndrome or a stimulant exposure changes what the cardiac team monitors and treats.

The third is to separate acute distress proportionate to a frightening event from a psychiatric disorder requiring treatment. The responses differ, and conflating them either medicalises normal fear or misses treatable illness.

The fourth is explanation, and it is underrated. The patient has been told they had something resembling a heart attack and also that their arteries are clean, and those messages conflict unless somebody supplies the mechanism reconciling them. A coherent, non-technical account of a surge, a stunned muscle and a recovering ventricle does real therapeutic work, replacing an incoherent threat with a comprehensible one.

The fifth is treatment and follow-up. **Where depression or anxiety is present, it is treatable in its own right, and a combined psycho-cardiologic rehabilitation approach is the model these patients might benefit from.** The obstacle is structural: a patient discharged with a normal angiogram is likely to leave with follow-up from neither service, and arranging it before discharge is the most useful line in the consultation note. Choice of agent in cardiac disease is a psychopharmacology matter and is handed over. Where consent arises, the operating principle is that capacity is decision-specific and time-specific, and the doctrine belongs to the mental health legislation notes.

RULE **Write the timed antecedent narrative on the day you are called.** Everything else in the consultation can be reconstructed later from records and from the patient. The hour-by-hour account of what happened before onset degrades within days, and it is the one datum that makes the case interpretable to anyone reading the notes afterwards.

IN INDIA

This build located no Indian series, registry entry, service description or cost datum for this syndrome, and none is supplied. What can be said structurally, without inventing a figure, is that recognition depends on early angiography, on repeat ventricular imaging to demonstrate recovery, and on a liaison service that receives the referral before discharge, and that the availability of each varies widely across Indian settings. A resident asked about the Indian picture is on firmer ground naming those dependencies and the absent data than quoting a prevalence figure with no Indian denominator.

GOOD TO REMEMBER

On the ward, hold these points. A clean angiogram excludes an obstruction and excludes nothing else. A psychological trigger never excludes a coronary occlusion. And the muscle is stunned rather than dead, which is why the story ends in recovery, and why the psychiatric work of characterising the antecedent, treating comorbid illness and arranging follow-up is done during an admission everyone else treats as resolved.

MNEMONIC

Surge, Switch, Stun, Settle

Surge is the catecholamine load from adrenal and neural sources. Switch is the apical receptor coupling change from stimulatory to inhibitory. Stun is living but non-contractile myocardium. Settle is the recovery that makes the diagnosis retrospective. The psychiatric work attaches at the start, on the antecedent, and at the end, on comorbidity and follow-up.

Beyond one territory	Stunned	Common	Timed narrative
THE WALL-MOTION FEATURE THAT MAKES OCCLUSION IMPLAUSIBLE	NOT INFARCTED, WHICH IS WHY IT RECOVERS	DEPRESSION AND ANXIETY IN THIS POPULATION, STATED QUALITATIVELY	THE DATUM ONLY THE PSYCHIATRIST CAN OBTAIN

A surge of catecholamines stuns living myocardium, the presentation copies an acute coronary syndrome exactly because the muscle really is injured, and the diagnosis is completed by recovery rather than by the angiogram. Take the timed antecedent on the day, look for exogenous and withdrawal catecholamine load, treat the comorbid depression and anxiety, and arrange the follow-up a normal angiogram will otherwise cancel.

Renal, transplant and perioperative psychiatry

11 · Pre-transplant psychiatric evaluation and post-transplant neuropsychiatry of immunosuppressants (steroids, tacrolimus, ciclosporin)

A transplant psychiatric opinion is read by a committee allocating a scarce organ, which makes it the one psychiatric assessment whose reasoning has to be legible to people outside

psychiatry. The topic splits cleanly and the sides are examined differently. Before transplantation, the question is what the psychosocial assessment must contain and how its conclusion should be written. Afterwards, it is which neuropsychiatric syndromes the immunosuppressants produce, and how to keep that possibility open while a recipient with several plausible causes of a new confusional state is worked up. The routing is strict, and one part of it is a gap rather than a destination. The general machinery by which a co-prescribed agent alters exposure, the inhibitor and inducer ladders and the cytochrome pathways, belongs to the psychopharmacology notes; that chapter names no immunosuppressant, so the interaction rankings and dose changes specific to these drugs are stated nowhere in this series and must not be reconstructed from memory. The confusional state itself belongs to the stroke, traumatic brain injury, delirium and focal brain syndromes notes. The steroid-induced psychiatric syndrome is routed on the plate below and taught in the endocrine and metabolic notes, while dialysis and the intensive-care sequelae around a transplant admission sit in the organ-interface routing section.

11.1 Where the opinion's form is owned, and what this section adds

Two limbs of this assessment are owned elsewhere and are named here rather than taught. How the opinion itself should be written, and the dual agency the interview runs under, belong to the Consultation-liaison, geriatric and transcultural psychiatry notes; so does the separate assessment of a living donor, where the question is whether the offer is voluntary and free of inducement, and where the psychiatrist works for the donor alone. The statutory framework governing donor authorisation belongs to the mental health and medical legislation notes. What this section adds is the limb those chapters do not carry: what the psychosocial assessment must contain, and how a capacity finding made during organ failure should be scoped, dated and re-asked.

11.2 The psychosocial assessment, and the limb this packet grounds

Guidance on evaluating candidates for kidney transplantation specifies the content of the psychosocial assessment, and the limb concerning psychiatry is stated explicitly. The elements of that assessment should include **a mental-status examination**, together with cognitive evaluation to ensure valid decision-making capacity and the ability to provide informed consent for transplantation.

Read that carefully, because the source is narrower than the way the topic is usually taught. It attaches the requirement to the psychosocial assessment as such, and attaches no timing qualifier of its own. This fragment therefore states it at that scope, and treats "pre-transplant" as a description of the setting in which the assessment is customarily performed rather than as part of what the guidance recommends.

The cognitive limb is the one candidates under-weight and the one most likely to change the answer. Organ failure impairs cognition by several routes at once, and the assessment often happens at the point of greatest impairment: uraemia, hepatic encephalopathy, hypoxia, poor cerebral perfusion and disrupted sleep. A patient examined in that state may appear unable to hold a decision they could hold comfortably once dialysis, oxygenation or ammonia handling

improved. So date the finding, name the reversible contributors, and set a point at which the question is asked again.

That licenses the one doctrinal point this section may state. Capacity is **decision-specific and time-specific**: a person may hold capacity for one decision and lack it for another on the same afternoon, and may regain it as their physiology improves. The statutory test, and the machinery that follows a finding of incapacity, are taught in the mental health legislation notes and named here rather than reproduced.

TRAP **Recording a capacity finding as a property of the patient rather than of a decision at a time.** In a population whose cognition tracks a treatable metabolic state hour by hour, writing "lacks capacity" as though it were a diagnosis removes a patient from consideration on a measurement taken at their worst, and the harm is invisible in the notes.

11.3 The rest of the interview, and where its evidence lives

Beyond the capacity limb, the assessment covers ground this packet does not ground with a source, so the honest move is to teach the structure of the interview rather than thresholds nobody here can cite. The domains a unit expects an opinion to address are the psychiatric history and current mental state; substance use and what any stated abstinence rests on; the adherence record; the caring arrangements after discharge; the patient's own account of the regimen and its permanence; and the financial arrangements that will keep immunosuppression running once the admission ends.

One interviewing discipline deserves stating outright, because it is where weak assessments go wrong. Ask about behaviour already on the record in preference to intention offered in the room. The record can be checked: dialysis attendance, missed appointments, prescription collection, what happened the last time a regimen became inconvenient. Why a stated intention is the weaker evidence here, given who the opinion is written for, is argued in the Consultation-liaison, geriatric and transcultural psychiatry notes. That is a point about where evidence lives, and this packet holds no figure licensing any claim that a given behaviour predicts a graft outcome.

Mental state and cognition	Record, not intention	Decision and time specific	Dated, then re-asked
WHAT THE ASSESSMENT MUST CONTAIN	WHERE ADHERENCE EVIDENCE LIVES	HOW CAPACITY IS SCOPED	WHAT A COGNITIVE FINDING IS

11.4 Tacrolimus, and the shape of calcineurin-inhibitor neurotoxicity

The calcineurin inhibitors are what make post-transplant care a psychiatric topic rather than a purely medical one. Their neurotoxicity is prominent enough in the labelling to belong on the differential of every neuropsychiatric presentation in a recipient. The tacrolimus prescribing information states the shape directly: tacrolimus may cause a spectrum of neurotoxicities, and the most severe of these include **posterior reversible encephalopathy syndrome (PRES), delirium, seizure and coma**.

Treat the word spectrum as the examinable part. A label describing a spectrum and then naming only its severe end is telling you a milder end exists while leaving it undefined. This packet holds

no row describing that milder end, so this fragment names the severe end the source names and goes no further.

What the statement licenses is a low threshold, and not a diagnosis. A recipient on tacrolimus who develops a seizure, a fluctuating confusional state or a depressed level of consciousness has a drug on the differential from the first minute, alongside infection, metabolic derangement, graft dysfunction and cerebrovascular events. The observation raising the drug's rank is temporal: whether the change followed an alteration in the regimen, a change of formulation, or the addition of another agent. Which agents do that, in which direction, and what dose response is then required are prescribing questions owned by the psychopharmacology notes.

Note the boundary inside that list: delirium appears as a manifestation of the drug's toxicity, which puts the drug into the delirium differential without putting delirium assessment into this section.

11.5 Ciclosporin, and a signal the label attaches to a combination

Ciclosporin is the other calcineurin inhibitor in routine use, and its labelling carries a neurotoxicity statement structured differently from the tacrolimus one, which is why it is worth learning separately. The statement is conditional on a co-prescription: convulsions have been reported in adult and paediatric patients receiving ciclosporin, particularly in combination with high-dose methylprednisolone.

Several features of that sentence repay attention.

- The signal is attached to a combination. A convulsion in the days after pulsed corticosteroid for suspected rejection is a combination event, and the history must establish what else was given and in what order, rather than settling on whichever agent the patient has taken longest.
- The labelling names adult and paediatric patients alike, so the statement spans the age range and a child or adolescent recipient attracts the same consideration.
- The wording reports occurrence rather than a rate. It licenses vigilance and a position on the differential; it licenses no statement about how often this happens.
- The named partner is a corticosteroid, so the same episode sits inside the corticosteroid material routed below. Naming the combination is this section's job; the steroid syndrome is not.

The plate below arranges the direct neuropsychiatric toxicity of the immunosuppressants by agent class rather than by syndrome, which is the arrangement answering the ward question: you know what the patient is taking and you want to know what it can do. The calcineurin inhibitors and the mTOR inhibitors are marked as interaction-sensitive, and that mark is a pointer rather than a teaching, since the interaction rankings and dose changes they oblige sit in the psychopharmacology notes. This packet holds no row describing an mTOR-inhibitor neuropsychiatric syndrome, so the prose asserts nothing about that class.

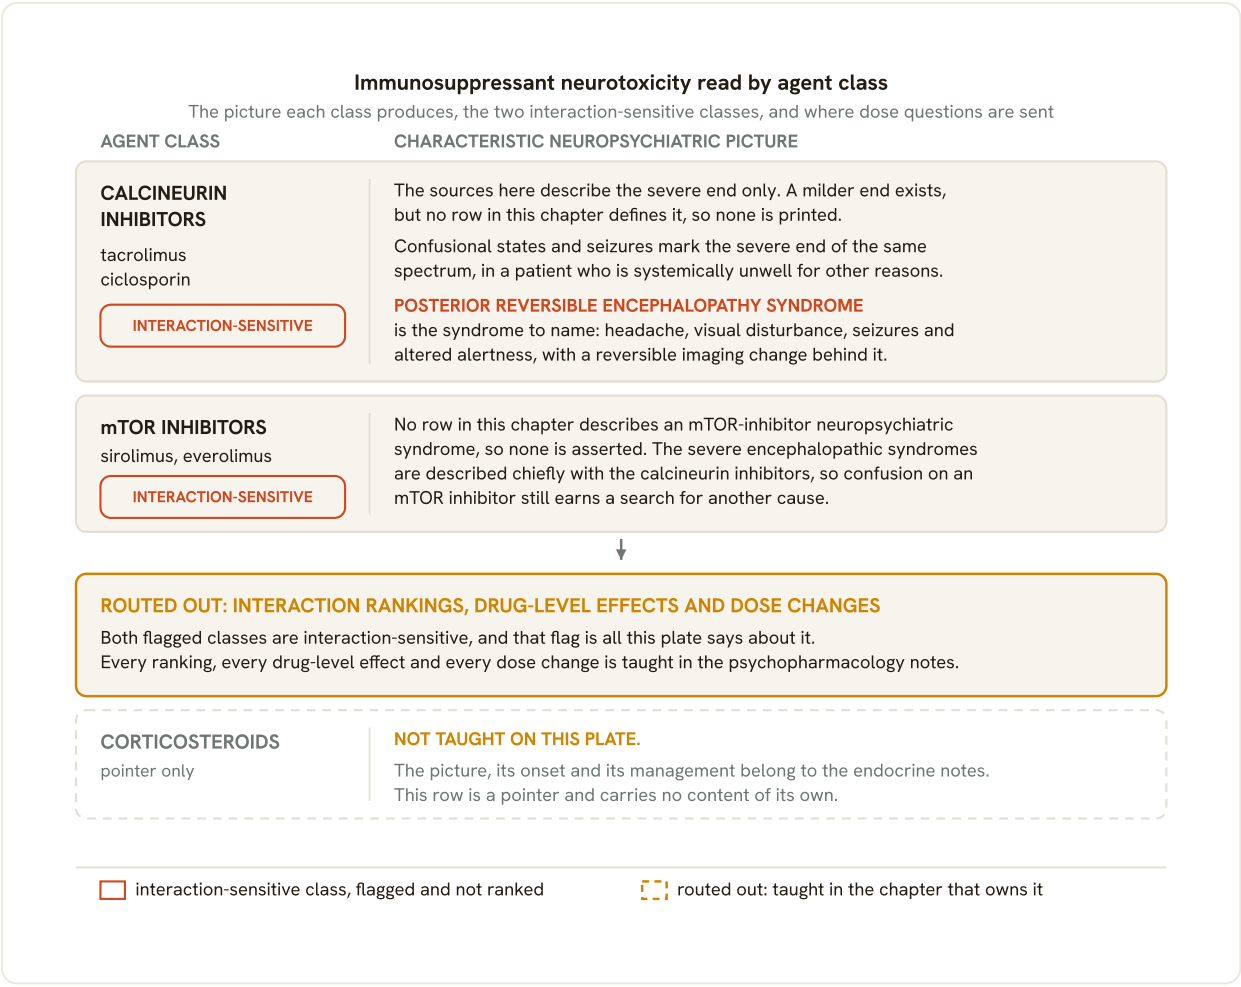


Fig 40.5 Direct neuropsychiatric toxicity of the immunosuppressants, read by agent class rather than by symptom. The calcineurin inhibitors carry the fullest picture: confusional states and seizures at the severe end, and posterior reversible encephalopathy syndrome as a named syndrome rather than a loose description of a confused transplant recipient. The milder end of that spectrum is left undefined here, deliberately, because no source in this chapter's evidence defines it, and naming a spectrum while supplying only one end of it from memory is how a plate acquires content its text does not hold. Nothing is asserted about an mTOR-inhibitor neuropsychiatric syndrome either, for the same reason; what the plate does say is that the severe encephalopathic syndromes are described chiefly with the calcineurin inhibitors, so confusion in a patient taking an mTOR inhibitor still earns a search for another cause before the drug is blamed. Three limbs are deliberately absent from the plate. Interaction rankings, drug-level effects and every dose change are taught in the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes; the corticosteroid picture is taught in the Endocrine and metabolic psychiatry notes; and the form of the pre-transplant opinion, together with the separate living-donor interview, belongs to the Consultation-liaison, geriatric and transcultural psychiatry notes. What that chapter records as a declared gap, the contents of the psychosocial assessment, is taught here. No Indian incidence figure is printed here, because the record was not located. From memory you should be able to name the class, give its severe picture, say the one syndrome out loud, and say which limbs this plate deliberately does not carry.

11.6 Corticosteroids, named here and taught elsewhere

Corticosteroids belong to the transplant regimen and therefore to this chapter, but the psychiatric syndrome they produce is owned elsewhere, so this section states the toxicity at label level and routes the clinical construct. Behavioural and mood changes sit among the common adverse reactions listed for corticosteroids, and that is the level at which this fragment may state them.

Steroid-induced psychiatric disturbance itself, meaning the mania, depression and psychosis that follow corticosteroid exposure and the management of each, belongs to the organ-interface

routing section of this chapter and to the endocrine and metabolic notes. An examiner asking about it in a recipient is asking a question whose answer is the same as in any other steroid-exposed patient: the transplant setting adds a co-prescribed calcineurin inhibitor and an immunosuppressed host, and adds no separate steroid syndrome.

11.7 The recipient with new neuropsychiatric signs, and the order of the questions

The characteristic difficulty after transplantation is that several sufficient causes are present at once, so the examinable skill is sequencing rather than selection. An immunosuppressed patient may be septic while barely febrile, metabolically deranged from graft dysfunction, hyponatraemic, hypoxic, and established on a calcineurin inhibitor and a corticosteroid at once, while carrying the psychological load of an operation they may have waited years for. An answer that names the drug and stops has selected the interesting cause rather than the first one.

The sequence that survives scrutiny puts reversible and dangerous causes first, keeps the immunosuppressant permanently on the list rather than at the end of it, and treats the psychological formulation as concurrent rather than as a diagnosis of exclusion. A recipient's distress about the donor, the graft, or a life reorganised around a regimen is real whether or not a metabolic cause is found, and deferring it until everything else is excluded means it is never addressed.

WHAT YOU OBSERVE	WHAT IT OBLIGES YOU TO ASK	WHERE IT IS TAUGHT
Fluctuating confusion or altered arousal	What is the medical cause, and what is the current regimen	Stroke, traumatic brain injury, delirium and focal brain syndromes notes
Seizure on a calcineurin inhibitor	What else was given, in what order, and when relative to the event	Label level here; prescribing in the psychopharmacology notes
Mania, depression or psychosis after corticosteroid exposure	Which agent, over what course, and what is the earlier psychiatric history	Organ-interface routing section; endocrine and metabolic notes
Doubt about a listing decision or an ongoing consent	Which decision, at what time, and what is reversible now	The requirement here; the statutory test in the legislation notes
Low mood on dialysis while awaiting a graft	What is the renal context, and what treatment is safe in it	Organ-interface routing section, which owns dialysis

Read the middle column rather than the first. Observation-to-diagnosis pairs fail on a stem where the observation has several sufficient causes; observation-to-question survives it, because the question stays the same whichever cause is operating.

11.8 The Indian setting, and what the assessment has to ask in it

Indian practice shapes this assessment more than any refinement of technique. Where the graft comes from a living relative, the psychiatrist acquires a second person to assess, with a different question and a different loyalty, inside a family in which pressure on a potential donor may be entirely unspoken. And because maintenance immunosuppression is a lifelong recurring cost, the sustainability of the regimen is a clinical variable rather than an administrative one.

IN INDIA

Ask the funding question in plain words, at the assessment and again before discharge: who will pay for immunosuppression once the admission is forgotten, and what happens if the recipient was the household's earner. Ask the travel question the same way, since a review schedule assuming a nearby clinic quietly excludes families who cannot make the journey. Where the donor is a relative, hold the separate donor interview the Consultation-liaison, geriatric and transcultural psychiatry notes describe, and record what was asked about voluntariness rather than only the conclusion. This packet holds no Indian figure for transplant volumes, living-donor proportion, drug cost, or psychiatric morbidity in Indian cohorts, so this box states the questions and states no magnitude.

GOOD TO REMEMBER

A short set of sentences carries most of the marks. The psychosocial assessment should include a mental-status examination and cognitive evaluation of decision-making capacity and the ability to provide informed consent for transplantation. Tacrolimus produces a spectrum of neurotoxicity whose severe end includes PRES, delirium, seizure and coma. Convulsions have been reported with ciclosporin, particularly in combination with high-dose methylprednisolone, and behavioural and mood changes are listed among the common corticosteroid adverse reactions.

PITFALL

Attributing a post-transplant seizure to ciclosporin alone. The labelled signal is conditional on a combination with high-dose methylprednisolone, so an answer naming the calcineurin inhibitor and omitting the corticosteroid has reported part of the statement and has stopped taking the history where it was about to become informative. The reverse error costs the same marks.

The pre-transplant side this section owns is the content of the psychosocial assessment, a mental-status examination and a dated cognitive evaluation of capacity, with the opinion's own form owned elsewhere; the post-transplant side is a permanently open differential in which tacrolimus carries a spectrum reaching PRES, delirium, seizure and coma, ciclosporin carries a convulsion signal tied to high-dose methylprednisolone, and corticosteroids are named here and taught elsewhere.

12 · Perioperative psychiatric management, capacity for surgery and postoperative cognitive dysfunction

Surgery reliably produces a new psychiatric presentation in patients who had none the week before. Perioperative psychiatry is best learned as a sequence, because the questions arrive in a fixed order and each has a different owner. Before the operation the question is optimisation and measurement. At consent it is capacity. Afterwards it is which cognitive construct is being named. This section installs that sequence and makes the hand-offs a candidate is expected to make rather than answer: delirium assessment and management, taught with the stroke, brain injury and focal syndrome material, and capacity doctrine with its statutory test, taught with the mental health legislation material.

12.1 The perioperative sequence, and who owns each question

A psychiatrist meets a surgical patient in a small number of recognisable ways, and naming them is the fastest route to a structured answer. The first is the planned preoperative review of a patient with established psychiatric illness, asking whether the illness or its treatment changes the operation or the recovery. The second is the urgent review at consent, where the team is unsure whether the patient can make the decision in front of them. The third, by far the commonest, is the postoperative call about a patient who was lucid on admission and is now confused or agitated. The fourth, most often missed, is the late review months afterwards, where the complaint is not confusion but persisting difficulty with memory or work.

Those shapes belong to different literatures and different chapters, and an examiner wants them separated before the answer starts. The characteristic failure is a single undifferentiated account of postoperative confusion, which loses marks in both directions: it imports delirium teaching into a question about cognitive trajectory, and trajectory language into a question about a disturbed patient on the ward tonight. So fix ownership first. The acute confusional presentation is the delirium chapter's, entire; consent and its statutory machinery are the legislation chapter's; agent, dose and organ-adjusted prescribing are the psychopharmacology chapter's. The plate below traces the pathway from preoperative review, through the operation, to the postoperative neurocognitive outcome, marking each exit where it occurs.

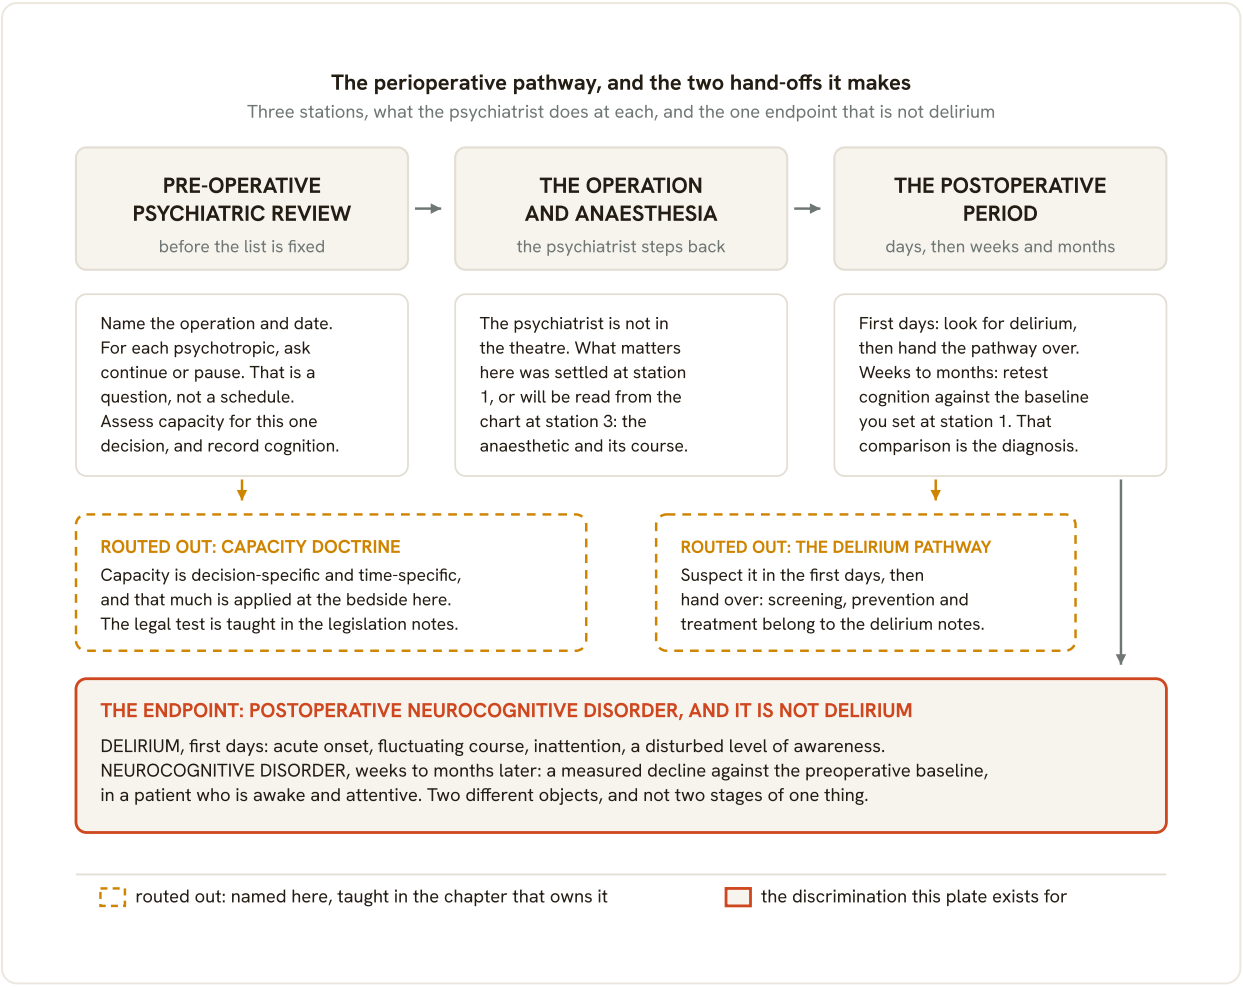


Fig 40.6 The perioperative pathway drawn left to right, with the psychiatrist's task at each station and the two hand-offs marked as branches off the rail. Station 1 is where the work is actually done: medication decisions are framed as questions rather than as a schedule, capacity is assessed for this one decision at this one time, and the cognitive baseline is recorded, because without that baseline the endpoint at station 3 cannot be measured at all. The capacity branch routes to the Mental health legislation and forensic psychiatry notes, which hold the legal test, and consent to the surgical decision is set out in the Forensic ethics, consent and legal procedure notes; the delirium branch routes to the Stroke, TBI, delirium and focal brain syndromes notes, which hold screening, prevention and treatment. The endpoint is postoperative neurocognitive disorder, and a nomenclature working group has recommended perioperative neurocognitive disorders as the overarching term for cognitive impairment identified in the preoperative or postoperative period. From memory you should be able to state what the psychiatrist does at each station, name both hand-offs, and separate delirium from postoperative neurocognitive disorder on onset, attention and course.

12.2 The name of the thing, and why the name was changed

The older vocabulary here was built by researchers rather than clinicians, and it showed. Terms for cognitive change after anaesthesia and surgery accumulated separately, each defined by whichever measurement a study happened to use, so a single patient could be described in incompatible ways depending on which literature the describer had read.

A nomenclature consensus working group addressed this by proposing a single umbrella. **The working group recommends that perioperative neurocognitive disorders be used as an overarching term for cognitive impairment identified in the preoperative or postoperative period.** Read that formulation slowly, because it does more than rename an old idea.

Notice what kind of statement it is: a recommendation about terminology, issued by a working group, and it should be presented as such. An answer promoting it into a diagnostic category with

criteria and thresholds has claimed more than the recommendation says. Notice next where the umbrella reaches. It covers impairment identified in the preoperative period as well as the postoperative period, which is the counter-intuitive limb and the examinable one. That decision converts preoperative cognitive assessment from a chore attached to the construct into part of the construct itself.

RULE **The umbrella reaches backwards, not only forwards.** The overarching term covers cognitive impairment identified before the operation as well as after it. A candidate who describes it as a class of postoperative complication has described half of it, and has lost the reason preoperative assessment belongs in the same discussion.

12.3 The preoperative psychiatric review, and what it is actually for

A preoperative review earns its place by settling a short list of questions where the surgical and anaesthetic teams will find the answers. The first is continuity of psychotropic treatment across the fasting window and any period in which the patient cannot swallow. The review does not choose the agent, which belongs with every dose and interaction ranking to the psychopharmacology material, but it must state, for each agent, whether it continues, is withheld or is given by another route, and which team owns that decision. A plan recording only that the patient takes long-term medication has settled nothing.

The second is withdrawal. An elective admission removes access to alcohol, to nicotine and to anything else the patient used at home, at a predictable hour. A withdrawal syndrome declaring itself on a surgical ward is read, by people looking for a surgical explanation, as postoperative confusion. Establishing what the patient used before admission is preoperative work, and recording it where the postoperative team will see it is the same job.

The third is the illness features that change postoperative risk without changing the operation: patients who communicate pain poorly, who do not seek help when something is wrong, who become agitated in unfamiliar surroundings, or whose treatment history makes an interruption expensive. Each is a practical problem invisible to a clerking that asks about organs. Pre-transplant psychiatric evaluation shares this shape with a different content and a different threshold, and is taught by the transplant section of this chapter.

12.4 The preoperative cognitive baseline, and why it is not optional

The most useful contribution a preoperative review makes to the postoperative period is a measurement taken before the operation. **Older patients should have preoperative cognitive assessment using established screening or diagnostic tools.** The recommendation names a class of instrument rather than a particular one, and this section follows it exactly. Naming a specific instrument here would need a source that attaches a property to it, and none is held, so supplying one from recollection is the error this chapter exists to prevent.

The reasoning is worth stating, because it is what makes the recommendation defensible rather than merely official. Postoperative cognitive change is a change, and a change is a relation between two observations; one postoperative observation cannot establish it. Without a

preoperative measure the team compares the patient against an impression of the patient, formed by relatives under stress and reported after the event.

There is a quieter second yield. Impairment found at the preoperative assessment is informative before any operation happens: it marks a patient whose recovery needs more support, it bears on how the consent conversation is conducted, and it frequently identifies impairment nobody had recognised.

TRAP **Reporting a low postoperative score as a postoperative deficit.** Candidates do this because a score looks like evidence and because the postoperative score is the one they have. Without a preoperative measure it establishes the patient's current level and nothing about change, and a patient with unrecognised pre-existing impairment generates the same result. The examinable point is not the score, it is the comparison the score was supposed to enter.

12.5 Capacity for surgery, and the part of it that belongs here

Almost all of capacity doctrine sits elsewhere. The statutory test and its limbs, the presumption the law starts from, the safeguards and the procedure that follows a finding of incapacity are taught with the mental health legislation material, and none of it is restated here or cited by provision. One property does belong in this section, because the perioperative period is where it does most of its work. **Capacity is described as decision-specific and time-specific.**

Decision-specific means the question is never whether the patient has capacity as a standing attribute, but whether this patient can make this decision. A patient may be unable to manage a complex financial arrangement and entirely able to decide about an operation, and the reverse holds equally. An opinion recorded for one purpose therefore does not answer a different question asked later.

Time-specific means the answer attaches to a moment rather than to a person, and the perioperative period supplies an unusual density of moments in which it can move: pain, disturbed sleep, premedication and the hours immediately after anaesthesia all sit inside the window in which consent is sought. An opinion formed on the morning ward round is not transferable to a consent conversation held that evening. What the psychiatrist supplies is a description of the patient's functional abilities for the specified decision at the specified time, in the terms the legislation sets out; the decision itself belongs to the treating team.

PITFALL

Recording a global verdict. A note stating that the patient lacks capacity, with no decision named and no time attached, is unusable by the team receiving it and indefensible if it is later examined. It also travels: written once, it is quoted for months, across decisions it was never about. Name the decision, name the time, confine the opinion to both.

12.6 Postoperative neurocognitive disorder distinguished from delirium

The two constructs are distinguished by what kind of thing each one is, not by comparing symptom lists. Delirium names an acute disturbance of attention and awareness with a characteristically fluctuating course. It is a condition in its own right, with its own screening

instruments, precipitant list and management, all taught with the stroke, traumatic brain injury and focal brain syndrome material and none of it repeated here. Postoperative neurocognitive disorder names something structurally different: a change in cognitive performance relative to a preoperative baseline, established by comparison rather than by observation at the bedside.

That difference in kind produces every other difference an examiner will want. One is recognised by looking at the patient now; the other cannot be, because a single observation carries no baseline. One fluctuates, which is why repeated observation detects it; the other is a persisting shift, which comparison rather than repetition finds. One declares itself on the ward; the other declares itself after discharge, and reaches a clinician only if somebody thinks to ask.

THE QUESTION	DELIRIUM	POSTOPERATIVE NEUROCOGNITIVE DISORDER
What the term names	An acute disturbance of attention and awareness	A change in cognitive performance against a preoperative baseline
How it is established	Observation of the patient's current state	Comparison of a later measure against an earlier one
Characteristic course	Fluctuating, often worse at night	Persisting, without the same fluctuation
Where it declares itself	On the ward, to the clinical team	After discharge, to the patient and family
Who teaches assessment and management	The stroke, brain injury and focal syndrome material	This section for the construct; psychopharmacology for any agent

RULE **Delirium is observed; a perioperative neurocognitive disorder is computed.** One is a state visible in the room, and the whole of its assessment lives in the chapter that owns it. The other is a difference between two measurements, which is why a service that never takes the first measurement can never make the diagnosis, however carefully it examines the patient afterwards.

The recommendation governing the acute limb is stated here only as a boundary. **Postoperative delirium screening should use a validated delirium score.** Which score, how it is administered and what follows a positive result are the delirium chapter's, and a candidate who reproduces an instrument here has answered another chapter's question while believing they were thorough. The perioperative point is narrower: screening is systematic and instrumented rather than impressionistic, which is a statement about how a surgical service is organised.

The constructs are not alternatives to choose between, and the same patient may pass through both, in sequence, after one operation. Whether an episode of postoperative delirium predicts the later cognitive trajectory is a question no source held for this section answers, and it is recorded below as a gap rather than taught.

MNEMONIC

Measure before, ask at, name after

Measure cognition before the operation, because a change needs an earlier observation. Ask the capacity question at the decision and at the moment. Name the construct afterwards, because delirium and a perioperative neurocognitive disorder are different kinds of thing with different owners.

12.7 Answering the postoperative call without importing the wrong chapter

The postoperative call is where most candidates meet this material, and the discipline that rescues an answer is to establish the shape of the problem before naming it. Establish whether a preoperative cognitive measure exists, because that determines what can be concluded at all. Establish the timeline, separating a disturbance that arrived over hours from a difficulty the family has watched for weeks. Establish whether the disturbance is of arousal and attention now, or of performance over time. Establish what else on the chart could produce it, since a surgical patient carries pain, immobility, infection risk and an interrupted medication list at once.

A postoperative note worth reading later states each of the following.

- Whether a preoperative cognitive assessment exists, where it is recorded, and what it showed.
- The timeline of the change, expressed against the operation rather than against the referral.
- Which construct is being described, and which is being excluded.
- The medication position: what was withheld, what was resumed, what is outstanding.
- Who reviews the patient after discharge, since the trajectory question cannot be answered on the ward.

Routing follows the same discipline. Psychological sequelae of a critical-illness stay, and delirium at the end of life, are taught by the organ-interface section of this chapter. Every question about agent, dose and organ failure belongs to the psychopharmacology material. Naming the destination answers that part of a question completely.

Before it, and after it	Measured, not recalled	This decision, this moment	Observed, or computed
WHAT THE UMBRELLA TERM COVERS	WHAT A LATER SCORE IS COMPARED WITH	HOW THE CAPACITY OPINION IS SCOPED	WHICH CONSTRUCT IS BEING NAMED

Every tile scopes something rather than reporting it: the reach of the term, the comparison a later score must enter, the boundaries of a capacity opinion, and which of the two constructs is in play.

12.8 The Indian perioperative reality, and what it changes

The perioperative recommendations held here come from a European perioperative safety guide, and the nomenclature recommendation from an international consensus working group. No Indian perioperative cognitive-assessment standard, and no Indian figure for how often preoperative cognitive assessment is performed, is held by any source behind this section. That absence is stated rather than filled, and it should be stated in an answer too.

What transfers is clear enough to reason about. The preoperative cognitive baseline is among the cheapest recommendations in perioperative medicine: a brief structured assessment, a person to administer it, and a place in the record the postoperative team will find. The barrier in most Indian surgical services is not cost but ownership, because preoperative clerking belongs to the surgical and anaesthetic teams, and the cognitive item is the first dropped when the list is long. Language and literacy add a second consideration: an instrument administered in a language the patient does not use, or whose items assume schooling the patient did not have, produces a low score that is a finding about the instrument rather than the patient.

IN INDIA

The practical move for an Indian unit is to attach the cognitive item to the anaesthetic preoperative assessment, where older surgical patients are already seen and documented, rather than to build a separate psychiatric pathway that depends on a referral being made. The psychiatric contribution is then to specify which instrument the unit uses, in which languages, administered by whom and recorded where, and to accept that a brief assessment applied consistently beats an ideal one done occasionally. No Indian coverage or outcome figure for such a pathway is held here.

GOOD TO REMEMBER

When the surgical team telephones about a confused postoperative patient, the first question is not what is wrong but whether anybody measured them before. If yes, a genuine comparison is available and a trajectory question can be answered. If no, say so plainly in the note, manage the acute presentation through the chapter that owns it, and arrange the review after discharge.

Delirium is seen at the bedside and belongs to another chapter; a perioperative neurocognitive disorder is a comparison against a preoperative baseline, which is why the measurement taken before the operation is this section's whole practical yield.

Oncology, palliative and paraneoplastic psychiatry

13 · Psychiatric effects of corticosteroids, chemotherapy and interferon (steroid psychosis, chemobrain)

The examinable content in cancer-drug neuropsychiatry is not a list of side-effects but a mapping: which agent, which syndrome, and over what interval that syndrome appears and is looked for. A candidate who has memorised that cytotoxic chemotherapy is associated with cognitive complaints, and that interferon-alpha is associated with depression, has learned the two least discriminating facts available. Marks sit on the rest: where in the course a syndrome belongs, what else produces the same picture at that point, whether the complaint has an objective correlate, and what the evidence licenses in advance rather than in response. This section installs that mapping for two agent families and deliberately routes a third away.

13.1 The boundaries this section keeps

Three boundaries govern what follows. Each exists because the material sits in full elsewhere in these notes, and a second version written here would drift from the first without anyone noticing.

The corticosteroid limb of this topic is not taught here. Systemic corticosteroids and steroid-induced psychiatric disturbance are routed on the chapter's opening plate to the endocrine and metabolic psychiatry material, and the merged organ-interface routing section of these notes carries the clinical construct. This section names steroids and goes no further with them.

Organ-adjusted prescribing is elsewhere too. Doses, dose adjustments and interaction rankings between psychotropics and cytotoxic agents belong to the psychopharmacology notes, and no figure of that kind appears below. The commonest way a candidate loses control of this topic is to answer a question about syndrome recognition with a half-remembered interaction table.

The third boundary is interferon. **Interferon-alpha** is taught here. **Interferon-beta** is taught in the Neuropsychiatry of systemic and neurological disease notes. That is a boundary rather than an accident of filing, and a later subtopic gives the reason.

RULE **Name the agent, name the syndrome, name the interval.** An answer that pairs an agent with a syndrome and stops has given half of what a ward needs. The interval tells the clinician when to look and which explanations compete at that moment.

13.2 Cancer-related cognitive impairment, and why the patient's word is the wrong construct

Patients say **chemobrain**. The construct a clinician should hold is **cancer-related cognitive impairment**, and the change of name is not fastidiousness. A term that asserts its own cause stops the clinician looking, and the value of this topic lies in what is still worth looking for after the complaint has been taken seriously.

The presenting complaint is characteristically about effortful cognition rather than frank forgetting: holding several things in mind, sustaining attention through a task, retrieving an available word, working at the old speed. Patients describe the change against their own baseline rather than any norm, which is both why the complaint is genuine and why it is hard to demonstrate.

Note what the study design in the next subtopic concedes before a single result is read. Its comparison groups are patients with the same cancers, not healthy volunteers. A design that needs disease controls is admitting that the disease, its investigations, its surgery, its other treatments and the fatigue around all of them are candidate causes of the same complaint. The cytotoxic agent is one contributor inside a bundle, and the patient's word names the one easiest to blame.

That study measured cognition with self-reported and performance-based cognitive measures, and it is worth registering that it used both. A complaint and a test result are different observations: the first cannot tell you whether performance has changed, and the second cannot tell you whether the person's life has changed.

13.3 The longitudinal evidence, read for exactly what it licenses

The anchor for the time course here is a longitudinal study of patients receiving chemotherapy for two tumour types, each type carrying its own control group. **225 participants** were assessed. The breast-cancer stratum held **41 patients and 36 controls**; the lymphoma stratum held **73 patients and 75 controls**. The reported conclusion is that, compared with controls, patients with breast cancer and lymphoma showed persistent cognitive worsening still present at **1 year** after treatment.

Read the boundary of that sentence carefully, because it is where the marks are. It is a group-level comparison: it describes two treated groups against two disease-matched groups and predicts nothing about the individual in front of you. It covers two tumour types, so it says nothing about regimens used for other cancers. The strata are modest, and the lymphoma stratum is the larger, so generalise more slowly from the smaller.

TRAP **Converting an observation window into a prognosis.** A finding reported up to a stated point after treatment describes how far the looking reached, not where the problem ended. Candidates translate it into reassurance, that it settles within the year, or into catastrophe, that it is permanent. The evidence supports neither, only that worsening was still detectable at the end of the window examined.

The consequence of persistence is a scheduling consequence. If the complaint outlasts the treatment, the point at which oncology hands the patient back is the point at which nobody is reviewing cognition, and it is exactly then that the patient is trying to return to work. A service that reviews cognition only while chemotherapy runs has arranged to stop looking just before the problem becomes disabling.

13.4 Assessing the complaint without reaching for an instrument

The task is not to confirm the patient's attribution. It is to characterise the complaint, subtract everything treatable that produces the same picture, and decide whether anything acute or evolving sits underneath it. The order matters: the treatable contributors are common and the acute ones are dangerous, and both are missed by a clinician who has already accepted the diagnosis the patient offered.

THE QUESTION TO ASK	WHY IT DISCRIMINATES	WHERE THE ANSWER ROUTES
Did the difficulty predate the first cycle?	Separates a complaint attributed to treatment from one already present at diagnosis	Keeps the differential open instead of closing it on the drug
Is the picture stable across the day, or fluctuating with clouded attention?	Fluctuating inattention is an acute-state question, not a chronic-impairment one	The stroke, traumatic brain injury, delirium and focal brain syndromes notes
Is mood, sleep, pain or fatigue carrying the complaint?	Effortful cognition fails when drive and sleep fail, and those are addressable now	The depression, demoralisation and adjustment material in the merged organ-interface routing section
Is the picture static, or evolving between reviews?	An evolving encephalopathy is a different class of problem from a stable deficit	The paraneoplastic neuropsychiatric syndromes section; the workup sits in the neurology-psychiatry interface notes
Does the complaint have an objective correlate on testing?	Self-report and performance are separate observations answering different questions	Decides whether the conversation is about workplace accommodation or further assessment

Two further routings belong in the same breath. Capacity, where a treatment decision turns on it, is decision-specific and time-specific, and the statutory test is taught in the mental health legislation notes. Where the complaint sits alongside a request to stop treatment, the palliative and existential material in the merged organ-interface routing section owns that conversation.

GOOD TO REMEMBER

The most useful thing a liaison psychiatrist adds here is a timeline. Ask when the difficulty began relative to diagnosis, to surgery, to the first cycle and to the end of treatment, and ask which task first went wrong. A complaint that predates any cytotoxic agent, one that fluctuates within a day, and one still worsening after treatment finished are three different problems wearing the same words.

13.5 Interferon-alpha, and the one thing this section can state as evidence

A spelling point first, because it costs marks in reading rather than in writing. The oncological preparation is written **interferon alfa** in drug nomenclature and in much of the trial literature, and interferon-alpha in general prose. They are the same agent.

The evidence this section can state concerns prevention rather than treatment. In patients with malignant melanoma who were eligible for high-dose interferon alfa therapy, pretreatment with paroxetine appears to be an effective strategy for minimising depression induced by interferon alfa. That was tested in a randomised comparison in which **40 patients were randomised, 20 to paroxetine and 20 to placebo**, with symptoms consistent with major depression during interferon alfa therapy as the outcome of interest.

Retain the qualification exactly as the source states it. Pretreatment appears to be an effective strategy, which is weaker than a claim that it prevents the syndrome, and an answer that upgrades

it has misreported the source. The population is equally load-bearing: melanoma, high-dose interferon alfa, one antidepressant, and pretreatment given before the interferon rather than treatment started once mood has changed. None of those boundaries transfers on its own to another cancer, another interferon, another antidepressant or a reactive prescription.

One thing can be inferred honestly from the design, and it carries the time course. A placebo-controlled prevention trial is only designable if the syndrome it prevents is expected often enough, and soon enough after therapy starts, to be counted inside the treatment course. The trial's shape therefore tells you that interferon-alfa depression is foreseeable and time-linked rather than idiosyncratic. That is an inference about the design, and this section publishes no rate.

RULE **Prophylaxis is a different claim from treatment, and a population is part of a claim.** The bound evidence supports pre-emptive antidepressant cover in a specified oncology population starting a specified interferon regimen. It says nothing about what to do once the mood syndrome has declared itself, and it does not extend to everyone about to receive an interferon. Quoting a prevention result as a treatment result is an invisible error unless you state the population every time you state the finding.

13.6 Interferon-alpha is not interferon-beta, and the separation is clinical

These are distinct preparations, given to distinct populations for distinct reasons, and their psychiatric literatures are separate. Interferon-alpha reaches the psychiatrist through oncology and through infectious disease, in courses that are intensive and time-limited. Interferon-beta reaches the psychiatrist through a chronic neurological illness carrying its own mood burden independently of any drug, which is why the Neuropsychiatry of systemic and neurological disease notes hold it and this section does not.

The separation matters for three reasons an examiner can ask about directly. First, the background rate of psychiatric disorder differs between the populations, so a figure pooled across both would describe no real patient and mislead in both directions at once. Second, the confounders differ: in the alpha setting the competing explanation is the cancer or the infection and its treatment, in the beta setting the neurological illness itself. Third, the clinical decision differs. In the alpha setting the live question is whether to cover a foreseeable, time-linked mood syndrome before a defined course begins; in the beta setting it is embedded in long-term management of a chronic disease, and taken up where that disease is taught.

TRAP **Answering "the neuropsychiatric effects of interferon" as a single topic.** A stem that says interferon without saying which preparation is testing whether the candidate knows the distinction exists. The answer that scores separates alpha from beta, states which one it will discuss, and says why pooling them would be uninformative. The answer that fails treats the word as one drug and produces a paragraph true of nobody.

13.7 Agent to syndrome to time course, which is what the plate carries

The plate below is the compressed form of everything above, built on three columns rather than two. The first names the agent family. The second names the syndrome, in the vocabulary a

clinician would use in a note. The third, which does the work, places that syndrome on the treatment timeline: before treatment, during the course, at handover, and after treatment has finished. Read the third column first when revising, because it converts a list of associations into something usable at the bedside.

Two features of the plate are deliberate rather than accidental. The corticosteroid limb is absent, because steroids are routed on the chapter's opening plate to the material that owns them. And interferon-alpha and interferon-beta are drawn apart rather than sharing a row, for the reasons given above.

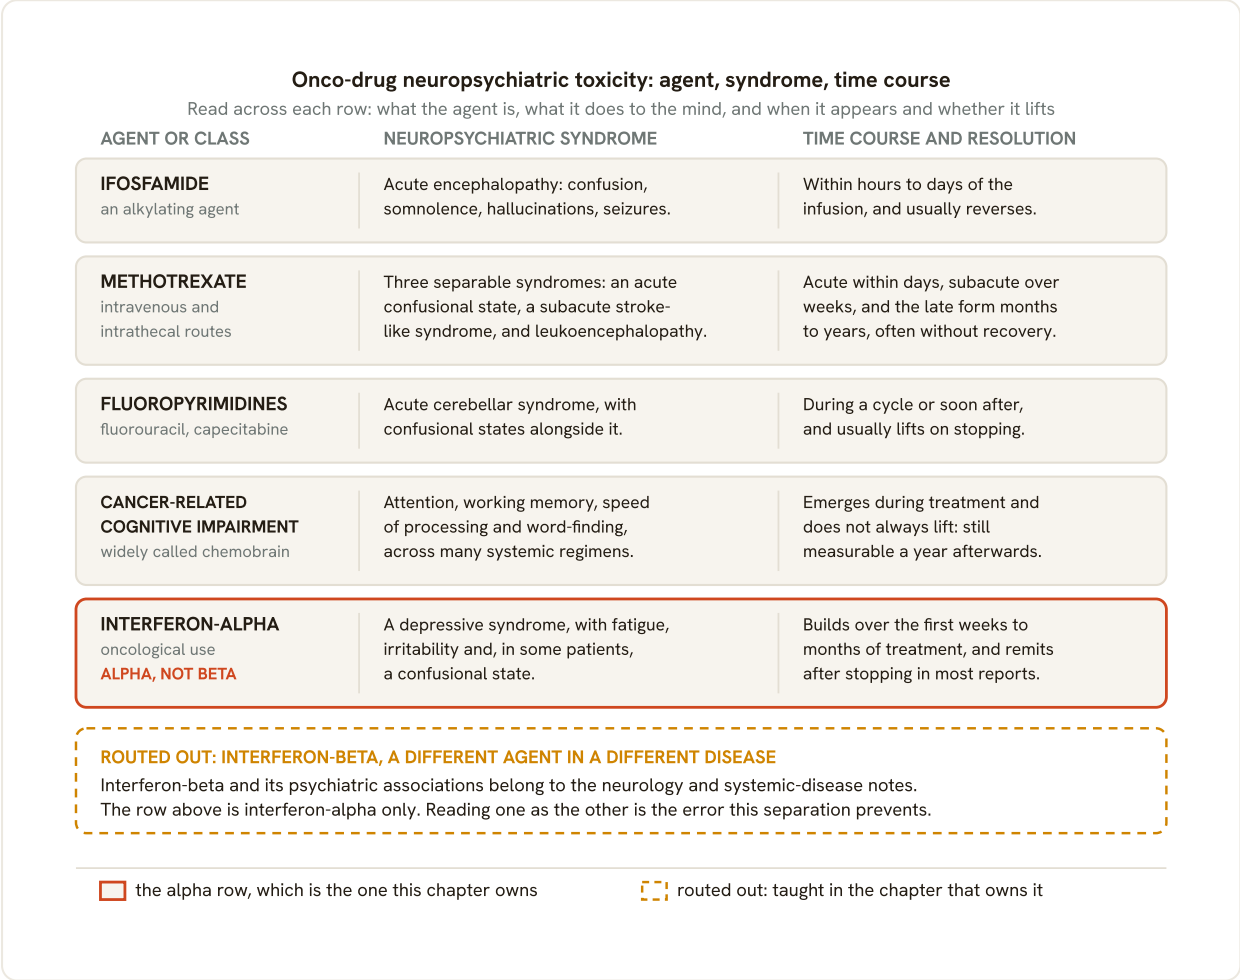


Fig 40.7 Chemotherapy and interferon neuropsychiatric toxicity read as agent, then syndrome, then time course, because at the bedside it is the time course that usually settles the differential. The acute encephalopathies declare themselves within hours to days and usually reverse on withdrawal, which is why ifosfamide, the fluoropyrimidines and the acute methotrexate syndrome sit together; methotrexate is the agent that also carries a subacute stroke-like syndrome and a late leukoencephalopathy that may never recover. Cancer-related cognitive impairment, widely called chemobrain, is the slow limb: it emerges during treatment, spares alertness, and has been measured as still present a year after treatment ends. Interferon-alpha produces a depressive syndrome that builds over the first weeks to months of treatment. The plate separates interferon-alpha from interferon-beta deliberately, because they are different agents used in different diseases and beta belongs to the Neuropsychiatry of systemic and neurological disease notes; the third drug group named in this topic is carried by the routing plate earlier in this chapter and by the Endocrine and metabolic psychiatry notes, so it is not drawn here at all. From memory you should be able to place any of these agents on the three columns and say whether the syndrome is expected to lift.

IN INDIA

No Indian figure for the frequency of cancer-related cognitive impairment, and none for interferon-alfa depression in Indian oncology practice, was located for this section, and none is asserted. What transfers is the structure. Where oncology and psychiatry sit in different institutions and follow-up is episodic, the practical question is not which cognitive instrument to buy but who reviews the patient once the oncology course closes, and when. Since the bound evidence describes a complaint still present after treatment ended, a service with no review point after discharge from oncology has designed itself to miss the window that matters most.

13.8 The shape of an answer that scores

A written answer here is graded on discrimination rather than length, and the discriminating content is ordered.

- Separate the agent families at the outset. Name the corticosteroid limb only to route it, and separate interferon-alpha from interferon-beta before saying anything about either.
- Use the construct rather than the patient's word, and say why the patient's word closes the differential prematurely.
- State every finding with its population attached. A finding without its population is not a shorter answer, it is a weaker and different one.
- Place each syndrome on the treatment timeline, and keep the observation window distinct from the prognosis.
- Route the neighbouring problems by name: delirium, paraneoplastic syndromes, depression and adjustment in cancer, capacity, and organ-adjusted prescribing.
- Say what is not known. The honest gaps here are large, and a candidate who names them is more convincing than one who fills them.

PITFALL

Treating the cognitive complaint as the end of the assessment because the patient has a plausible reason for it. A patient receiving chemotherapy has an explanation available for almost any neuropsychiatric presentation, and that availability is what makes delirium, an evolving encephalopathy and an untreated depression easy to miss. The plausible cause is the one to test last, not the one to accept first.

Agent, syndrome, interval
THE THREE THINGS AN ANSWER MUST NAME

Cancer-related cognitive impairment
THE CONSTRUCT THAT REPLACES THE PATIENT'S WORD

Prevention, not treatment
WHAT THE PRETREATMENT FINDING LICENSES

Alpha is not beta
THE SEPARATION A ONE-WORD STEM IS TESTING

Interferon-alpha is separated from interferon-beta and corticosteroids are routed away; the cognitive syndrome is cancer-related, not chemotherapy-proven; and every finding is stated with its population and its observation window attached.

14 · Paraneoplastic neuropsychiatric syndromes and onconeural antibodies: when to suspect occult malignancy

A paraneoplastic neuropsychiatric syndrome is a psychiatric presentation driven by the immune response to a tumour that has not yet been found. The final clause carries the whole examinable weight. This is not a complication of a cancer already sitting in the case notes; it is the event that makes someone go looking for the cancer in the first place. This section installs the reasoning that converts an ordinary-looking psychiatric presentation into a directed search for occult malignancy: what the term **paraneoplastic** actually claims, why the psychiatric clinic meets these patients ahead of the oncology clinic, what raises suspicion, what a psychiatrist owns and what is handed on, and which antibody has been reported to point where.

14.1 What the term claims, and the sequence that defines it

The prefix does the work. **Para** plus **neoplastic** names a disorder that accompanies a neoplasm while being made of something else. The damaged tissue is nervous tissue, and what damages it is an immune response the tumour provoked. The word **onconeural** names why such a response reaches the brain at all: the antigen belongs to the tumour and to neural tissue alike, so an attack aimed at one compartment lands in the other. Everything examinable here follows from that idea, including why the syndrome can be severe while the tumour stays small enough to have escaped detection.

The defining feature, for examination purposes, is a sequence rather than a severity. **The neurological or psychiatric syndrome comes first, and the malignancy is detected afterwards.** The task-force report that grounds this section states it plainly: **paraneoplastic neurological syndromes almost invariably predate detection of the malignancy**. Hold that sentence in mind through everything that follows, because it is what separates this topic from every other cancer topic on the day, and it is the reason the topic belongs to a psychiatry paper at all.

It is worth being precise about what the label has to displace before it is earned, because a neuropsychiatric picture in someone who turns out to have cancer has several ordinary explanations. It may be a direct effect of tumour mass or of secondary deposits. It may be metabolic or endocrine derangement produced by the tumour. It may be treatment: the cytotoxic agents, the interferon and the corticosteroids, each taught in its own place, in the sibling section on chemotherapy and interferon effects and in the routing section that owns steroid-induced disturbance. It may be adjustment, demoralisation or depression in someone absorbing a serious illness, which the routing section on cancer owns. Cross-refer each by name; this section restates none of their content.

Notice the asymmetry the sequence creates. In the ordinary cancer psychiatry of a known malignancy, those rival explanations are all available on the ward round, because the tumour, the metabolic panel and the drug chart are already in front of you. In the paraneoplastic case they are unavailable at presentation, because nobody has yet established that there is a tumour. The

clinician holds a syndrome and a history, and the question of malignancy has to be raised rather than handed over.

RULE **The defining feature is the sequence, not the severity.** **A paraneoplastic syndrome characteristically declares itself before the malignancy is detected**, which is why the phrase in the bank item is "when to suspect occult malignancy" and why an answer that describes a florid syndrome in a known cancer patient has answered a different question.

14.2 Why the psychiatric clinic meets these patients first

There are honest structural reasons why this presentation lands with a psychiatrist, and a candidate who can state them is well placed to answer the item. The first is the sequence itself. Because **the syndrome usually arrives before the malignancy has been detected**, the first clinician to see the patient is whichever clinician the presenting complaint routes them to. Where that complaint is behavioural, meaning altered mood, altered personality, psychosis, failing memory, reversed sleep, agitation or apathy, that clinician is a psychiatrist. The referral letter says nothing about cancer, because at that point there is nothing about cancer to say.

The second reason is that psychiatric formulations are generous, and their generosity is usually a strength. A subacute behavioural change in an adult can be absorbed comfortably into a first episode of psychosis, a late-onset depressive episode, a mood disorder cycling faster than usual, a dissociative presentation or a functional cognitive disorder. Each is internally coherent and each generates a plausible treatment plan. None of them obliges the clinician to ask what else could be producing this picture, so the question of occult malignancy is one the reasoning process will never raise unless the clinician raises it deliberately.

The third reason is what the examination is actually testing. The bank item asks when to suspect, and suspicion is a skill with a location in time: it belongs at first contact, alongside the initial formulation, and it loses value the longer it is deferred. A suspicion raised after several failed treatment trials is a lesser clinical act than the same suspicion raised at intake, because in the interval the syndrome has progressed while **the malignancy has continued to sit undetected**.

TRAP **Treating the difficulty as a laboratory problem.** Candidates picture the failure as a missed antibody result, so they write about panels and assays. The failure that actually happens is upstream and cheaper to fix: the question of an underlying tumour was never opened, so nothing was ever sent. An answer that begins with the test has skipped the step the item is examining.

14.3 Reading the presentation: tempo, company and fit

Suspicion is a discipline rather than a mood, and it should not be taught as vigilance in general. It is built from a small set of readings that a competent psychiatric history already produces, and the work is to read them deliberately rather than to hope they announce themselves.

The first reading is **tempo**. Psychiatric illnesses have characteristic time signatures, and a clinician who has taken many histories knows them implicitly. What matters here is the mismatch: a presentation that moves faster than the illness it resembles, that keeps moving while treatment is being delivered adequately, and whose trajectory is a slope rather than an episode.

Tempo costs nothing to record, requiring only a careful timeline with collateral corroboration, and it is the reading most often left vague in a psychiatric clerking.

The second reading is **company**, meaning what the psychiatric syndrome has arrived with. Psychiatric syndromes keep characteristic company, and neurological company is different company. When a behavioural presentation is accompanied by features belonging to the nervous system rather than to the mind, the formulation has acquired something it cannot metabolise, and that is informative in itself. The plate below sets out the clusters that carry this weight, and it is the right place to learn them, because they are the content of the source criteria rather than of this author's recollection.

The third reading is **fit**: whether this person fits the illness being diagnosed, and whether a general history has been taken at all. Psychiatric clerking has a known habit of collecting an exhaustive developmental history alongside a thin systemic one, so constitutional features, general medical risk and the ordinary questions an internist would ask are frequently missing. A history that never asked them cannot be reassuring, and the difference between a negative finding and an unasked question is one this chapter returns to repeatedly.

The plate that follows carries the working set: the onconeural antibody group, the tempo and sign clusters that raise suspicion, and the boundary against the autoimmune-encephalitis workup taught elsewhere. Read the plate as the operational list, and this prose as the reasoning that tells you when to reach for it.

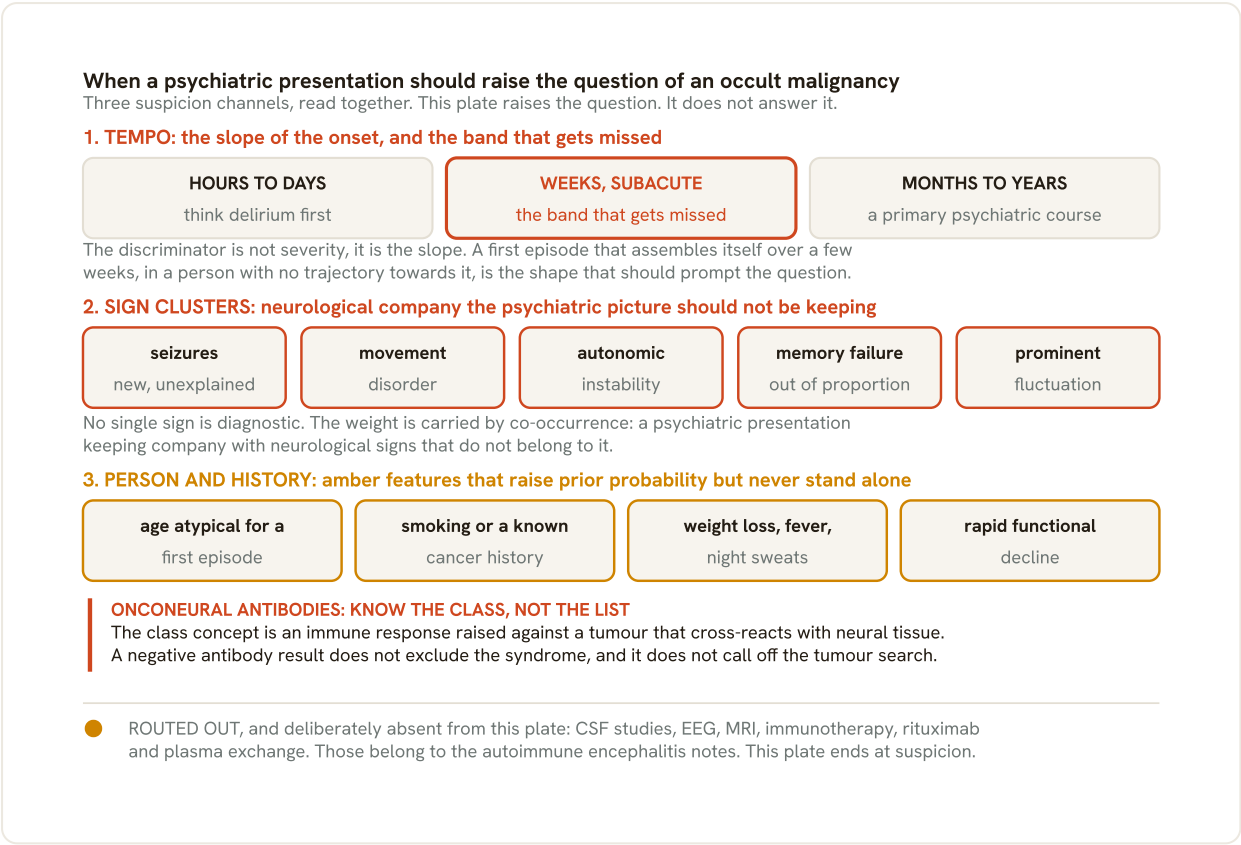


Fig 40.8 The three suspicion channels for an occult malignancy presenting psychiatrically. The plate is deliberately a suspicion instrument and stops there, because the workup and its immunotherapies belong to the autoimmune encephalitis notes and are cross-referenced out rather than duplicated. What it adds to the prose is the tempo rail, which puts the subacute band between two safer neighbours and shows why that band is the one that gets missed: a presentation over hours to days is read as delirium and gets investigated, a presentation over months to years is read as primary psychiatric and gets followed, and the few-week presentation falls between two reflexes. The onconeural antibody idea is carried as a class rather than a list because that is how it is examined and how it behaves, and because a negative result narrows nothing. From memory you should be able to state the tempo band that should worry you, name the sign clusters that a psychiatric picture should not be keeping company with, and say what a negative antibody result does and does not permit you to conclude.

14.4 The boundary: suspicion is yours, the workup is not

This section owns paraneoplastic suspicion. It does not own the autoimmune-encephalitis workup. Cerebrospinal fluid examination, electroencephalography, magnetic resonance imaging, first-line immunotherapy, rituximab, plasma exchange and intravenous immunoglobulin are taught in the Neuropsychiatry of systemic and neurological disease notes, and this text restates no parameter of any of them. The boundary is named explicitly because it is the one an author is most tempted to cross: the workup has protocols and sequences and feels like the substance of the topic, and it is the substance of a different topic with a different owner.

The split is clinical rather than merely editorial. The investigations and the immunotherapy form one coherent body of knowledge, learned and examined together. Suspicion is a separate skill, exercised by a different clinician at an earlier moment, and it is the skill this bank item names.

What the psychiatrist owns is worth stating positively, and it is short enough to hold as a list.

- The timeline, recorded precisely enough that tempo can be read off it later by a clinician who was not in the room.

- The collateral history, often the only source of that tempo.
- The general and systemic history that psychiatric clerking tends to abbreviate.
- The referral itself: one that says "for neurology opinion" transmits only the patient, while one that states the tempo, the company and the explicit reason a tumour is being considered transmits the reasoning.

THE QUESTION IN FRONT OF YOU	WHERE IT IS ANSWERED
Could this psychiatric presentation be driven by a tumour nobody has found yet?	This section.
What do the cerebrospinal fluid, the electroencephalogram and the scan show, and what immunotherapy follows?	Neurology-psychiatry interface notes, autoimmune encephalitis. Never taught here.
Is this delirium, and how is it assessed and managed?	Stroke, traumatic brain injury, delirium and focal brain syndromes notes.
Does this patient have capacity, and what is the statutory test?	Mental health legislation notes. Capacity is decision-specific and time-specific; the test is taught there.
How is the psychotropic adjusted for hepatic or renal impairment?	Psychopharmacology notes. This chapter adds no further version of a dose adjustment.
Is the disturbance caused by the cytotoxic agent, by interferon, or by the corticosteroid?	The sibling section on chemotherapy and interferon effects, and the routing section owning steroid-induced disturbance.
Is this neuropsychiatric lupus, neurosarcoidosis, central nervous system vasculitis or antiphospholipid syndrome?	The sibling sections owning each of those.
Is this depression, demoralisation or adjustment in someone already known to have cancer?	The routing section on cancer.

14.5 The screening imperative, and the guides it names

Once suspicion exists, the criteria are direct about what follows. **Cancer screening should be promptly undertaken when a paraneoplastic neurologic syndrome is suspected, and should be guided by the type of phenotype or antibody.** That single sentence is the operational core of this section, and it repays reading in parts.

The first operative word is **promptly**. It fixes a position in the sequence: the screen belongs at the point of suspicion rather than at the point of diagnostic certainty. The claim carries urgency and carries no interval, so this text supplies none, and a candidate who wants a period in days or weeks should take it from the criteria rather than from a note that would be inventing it.

The second operative phrase is **guided by the type of phenotype or antibody**. The instinct under pressure is to describe a broad sweep, imaging everything and hoping something appears. The criteria describe a search with a direction, where the clinical phenotype narrows the field of plausible primaries and the antibody, when available, narrows it further. An answer that reproduces the urgency while dropping the guidance has reproduced the easier part.

One inference is worth stating, because it is reasoning rather than recollection. If **the syndrome characteristically presents before the malignancy is detected**, the screen is being run in a body where the tumour has so far evaded detection. A first look that finds nothing has therefore not settled the question the way a first look normally settles one, and the posture that follows is continued attention rather than closure.

RULE **Prompt and guided, in that order and with equal weight. The criteria couple the timing of cancer screening to the suspicion of a paraneoplastic neurologic syndrome, and couple its direction to the phenotype or the antibody.** Reproduce the pairing. An answer carrying only urgency describes panic; an answer carrying only direction describes a screen that may arrive too late to matter.

Now the honest limit of this section, narrowed to what the ledger actually withholds. The antibody pairings are cited and taught in the subtopic below, so they have left the list of absences. What is still missing is the machinery around the screen: no row names which imaging modality corresponds to which phenotype, no row gives an interval for repeating a screen that found nothing, and no row states how long surveillance continues. Those three are what a candidate will want next, and they are owed to the source criteria rather than to this note.

14.6 Onconeural antibodies: the risk tier, the phenotype and the reported pairings

The term **onconeural antibody** names an antibody directed against an antigen shared by tumour tissue and nervous tissue. That definition is why the antibody is diagnostically useful at all: because the antigen belongs to each compartment, a serological finding points at a body region that has yet to be searched. Take the term seriously as vocabulary, because it carries the logic of the whole section in a single word.

The antibody earns its place here for the reason the criteria give: **the cancer screen is guided by the type of phenotype or antibody**. The antibody is a laboratory result while the syndrome is a clinical entity, so the result informs the search rather than constituting the diagnosis. And it sits alongside the phenotype in the claim as written rather than replacing it, so a clinician who has a phenotype and awaits a result already holds one of the guides the criteria name.

The criteria then sort the antibodies by the cancer risk each one carries, and that sorting is what makes the guidance operational. **The first group holds the high-risk antibodies, defined as those associated with an underlying cancer in more than 70 per cent of the patients who carry them.** A structural fact travels with the group: **most high-risk antibodies target intracellular antigens**. Put the two together and an inference follows. An antibody raised against an antigen sitting inside the cell is poorly placed to be doing the damage by binding it, so the finding is best used as a pointer towards a hidden tumour rather than as an account of mechanism.

One phenotype belongs to psychiatry more than to any other specialty. **Paraneoplastic limbic encephalitis is characterised by personality change, irritability, depression, seizures, memory loss and sometimes dementia.** Read that list as a psychiatric referral rather than a neurological one. Personality change, irritability and depression are what a psychiatric service is

built to receive; the seizure and the memory failure are what eventually move the formulation, often long after a psychiatric label has been applied and acted upon.

The reported pairings follow, with the evidence grade placed second so that it is read before the pairing it qualifies, and with the two rows that rest on more than one patient placed first. Read the second column before the third. A single reported case establishes that a pairing has been observed in a human being, and nothing beyond it: no frequency, no claim that this tumour is the one the antibody usually points to, and no entitlement to the word "classical".

ANTIBODY	WHAT THE CITED ROW IS	SYNDROME REPORTED WITH IT	TUMOUR REPORTED WITH IT
CRMP5	**The commonest associated diagnoses in a 24-patient cohort.**	**Peripheral neuropathy and cerebellar ataxia**	**Lung cancer**
Anti-Ma2	**Described as a specific marker for that pairing.**	**Limbic and brainstem encephalitis**	**Testicular germ-cell tumour, mainly in adult males**
Anti-Hu	**A single reported case.**	**Adie's pupil with sensorimotor polyneuropathy**	**Primary mediastinal small-cell carcinoma**
Anti-Yo	**A single reported case, arising years after the tumour was resected.**	**Cerebellar degeneration**	**Ovarian cancer**
Anti-Ri	**A single reported case.**	**Opsoclonus-myoclonus**	**Breast carcinoma, in a male patient**
Anti-amphiphysin	**A single reported case.**	**Stiff-limb syndrome**	**Invasive ductal carcinoma of the breast**

The second column is where marks are won and lost, which is why it now sits where the eye reaches it first. Four of these six rows rest on one patient each, and several of those reports were published precisely because they were unusual: an opsoclonus-myoclonus syndrome in a man, a cerebellar syndrome declaring itself years after the ovarian tumour was resected, a mediastinal rather than a pulmonary small-cell primary. A report selected for its rarity is the worst available basis for a statement about what is typical, and an answer that flattens this table into a list of classical associations has manufactured a confidence its own sources never carried.

The criteria also grade how much confidence the label deserves. **The PNS-Care score is a composite model for assigning diagnostic likelihood in a patient with a suspected paraneoplastic neurological syndrome**, and it answers in bands: **definite at 8 or above, probable at 6 to 7, possible at 4 to 5, and a score of 3 or less classed as something other than a paraneoplastic syndrome**. Learn the bands, and learn that a graded answer exists. The items the score weighs are absent from this fragment's ledger, so they are left to the criteria.

Syndrome, then tumour	Prompt and guided	Pointer, not diagnosis	Observed, not typical
THE SEQUENCE THAT DEFINES THE TERM	HOW THE CANCER SCREEN IS RUN	WHAT AN ONCONEURAL ANTIBODY IS FOR	WHAT A SINGLE-CASE PAIRING ROW ESTABLISHES

The four tiles carry the spine of the section: a sequence, a manner of screening, a use for the laboratory result, and a limit on what the pairing table licenses.

IN INDIA

The screening imperative assumes an onconeural antibody panel can be requested, paid for and reported within a useful period, and in much of Indian practice each assumption deserves checking before the plan is written. Panels are sent out to referral laboratories from most district and many teaching hospital settings, the cost is usually met out of pocket, and turnaround runs to the time a courier and a batch run take. The practical consequence is that **the phenotype is frequently the guide available first, with the antibody arriving later to refine the search**, which makes a carefully documented clinical picture more valuable here than where serology returns quickly. The Indian record on availability, pricing and reporting intervals for these panels was never located within this fragment's ledger, so it is described qualitatively rather than quantified.

GOOD TO REMEMBER

On the ward this reduces to a single habit. When a psychiatric presentation moves at the wrong speed, keeps company with neurological features, and sits in a person whose general history was never properly taken, write the timeline down and say the word "paraneoplastic" out loud in the referral. **The syndrome characteristically arrives before the malignancy is found**, so the psychiatrist is often the clinician best positioned to raise the question at the moment it is most useful. What follows, the fluid, the imaging, the electrophysiology and the immunotherapy, belongs to colleagues and to another section. Raising the question well is the part that is yours.

The syndrome almost always arrives before the malignancy is found, so the psychiatrist's part is suspicion at first contact: read tempo, company and fit, write the timeline down, say the word in the referral, and let the cancer screen be prompt and guided while the workup belongs to somebody else.

Autoimmune and inflammatory systemic disease, and burn injury

15 · Neuropsychiatric systemic lupus erythematosus: manifestations, diagnosis and management

The examinable difficulty in neuropsychiatric lupus is not the list of syndromes, it is deciding whether the syndrome in front of you belongs to the disease at all. Lupus reaches the nervous system often enough to produce almost any neurological or psychiatric presentation, and the same patient can also be septic, uraemic, hypertensive, taking drugs with central effects, and

frightened by an illness that has already taken things from them. This section installs the topic in the order the decisions are made: what the term covers, how the manifestations reach a psychiatrist, how attribution is reasoned when no test settles it, and what treatment looks like afterwards.

Several things this presentation touches are owned elsewhere and are named here rather than rebuilt. Delirium assessment and management, including the bedside instruments and the drug choice for an agitated delirious patient, belong to the stroke, traumatic brain injury, delirium and focal brain syndromes notes. Organ-adjusted prescribing belongs to the psychopharmacology notes, and the workup of autoimmune encephalitis to the neurology and psychiatry interface notes. Corticosteroid-induced psychiatric disturbance is taught by the routing section of this chapter that owns it and by the endocrine and metabolic notes, which matters here because the drug class used to treat this disease is itself a recognised cause of psychiatric disturbance. The antiphospholipid syndrome is owned by the sibling section that carries it alongside neurosarcoidosis and central nervous system vasculitis.

15.1 What the term is asked to cover, and why its breadth is the problem

Neuropsychiatric systemic lupus erythematosus, conventionally shortened to **NPSLE**, is a collective label for the neurological and psychiatric manifestations occurring in a person who has lupus. It is a container, not a diagnosis of the kind a psychiatrist is used to making. Standard usage sorts its contents into central nervous system events, peripheral nervous system events and psychiatric syndromes, and those groups share almost nothing at the level of mechanism, severity, investigation or treatment.

That breadth is the source of every difficulty that follows. A label covering both a mild headache and a catatonic psychosis cannot have one test, one mechanism or one treatment, and any sentence beginning "NPSLE is treated with" is suspect before its predicate arrives. The useful move is to handle the label as a question in parts: which syndrome, by which mechanism, in a patient whose disease is at what level of activity, with which alternative explanations excluded.

RULE **Name the syndrome, then the mechanism, then the attribution, then the treatment.** These are separate determinations and each can be right while the next is wrong. A candidate who says "this is neuropsychiatric lupus" has named a container. One who says "this is an acute confusional state, in a patient with active disease, in whom infection and metabolic causes have been excluded, treated as an inflammatory manifestation" can be examined on any of the four.

15.2 How the manifestations reach a psychiatrist

Grouping the manifestations by the route through which they arrive is more useful than grouping them by anatomy, because the route determines what you are being asked to do.

- **Diffuse presentations**, which psychiatry is most often called about: an acute confusional state, mood disturbance of either polarity, psychosis, cognitive impairment across attention, memory and executive function, and anxiety. Here the referral question is usually not "is this lupus" but "is this psychiatric".

- **Focal central events**: seizures, stroke-like episodes, movement disorder, myelopathy, aseptic meningitis, cranial neuropathy and a demyelinating syndrome. These reach neurology or medicine first, and psychiatry is called for the behavioural consequences, for capacity, or because the presentation was read as functional.
- **Peripheral events**: polyneuropathy, mononeuropathy, plexopathy, autonomic disturbance and a neuromuscular junction disorder. These rarely present to psychiatry directly, and their relevance is that pain and disability from them drive mood.

The classes are nomenclature rather than mechanism: they say where the disturbance sits and nothing about why. This chapter also holds no row giving the frequency of any manifestation in any population, so nothing here is called the commonest and no ordering by frequency is printed. Frequency statements circulate widely in this topic and disagree with one another, for the reason the next subtopic makes clear.

15.3 Attribution, which is the actual core of the topic

Everything listed above can be produced by lupus, and every one of them can be produced by something else in a person who happens to have lupus. Diagnosis is therefore **attribution**: a reasoned judgement that this syndrome, in this patient, at this time, belongs to the disease. It is reasoned rather than measured, because **there is no diagnostic gold standard by which the neuropsychiatric syndromes of lupus can be diagnosed**. That fact is the axis of the topic and belongs early in any answer.

Three consequences follow, each worth a sentence. Published frequency estimates disagree, because without a gold standard every study defines its own numerator: a series counting headache and mild mood symptoms and a series counting only severe syndromes are not measuring the same thing, and their percentages are not comparable even when both are correctly reported. No investigation is confirmatory, so every result is evidence to be weighed rather than an answer. And attribution is revisable, since a judgement made on admission may be overturned by the culture result, by the course, or by the response to treatment.

The reasoning runs on a small number of axes worth naming. The first is the temporal relation to disease activity elsewhere: a syndrome appearing while the disease is active in other organs is easier to attribute than one appearing in quiescent disease. The second is the presence or absence of a competing explanation, which the next subtopic takes up because it does the most work. The third is the character of the syndrome, since a course, phenomenology or examination finding atypical for a primary psychiatric disorder points towards a systemic cause. The fourth is response to treatment, and it is retrospective, which is why the attribution has to be made and acted on before the evidence for it is complete.

TRAP **Treating a positive antibody as an attribution.** Serology establishes that the patient has the disease. It does not establish that the syndrome in front of you was produced by it, and a patient with unambiguous lupus can have an unrelated psychiatric illness, an infection or a drug effect. Skipping the attribution step is also how a treatable cause is missed while immunosuppression is escalated.

15.4 What has to be excluded before the disease is credited

Because attribution rests so heavily on the absence of a competing explanation, the exclusion list is the substance of the assessment rather than preliminary housekeeping. Infection comes first and stays first in a patient immunosuppressed by both the disease and its treatment, and it includes meningitis, encephalitis, and the encephalopathy of systemic sepsis without central infection. Metabolic causes follow: uraemia, sodium and glucose disturbance, hepatic and thyroid disturbance. Hypertensive encephalopathy belongs in anyone with renal disease and a new confusional or seizure presentation.

Drug causes are the largest single trap here, because the treatment of the disease and the cause of the psychiatric syndrome can be the same agent. That construct is taught by the routing section of this chapter that owns it; what belongs here is the boundary. A new psychiatric syndrome beginning after a change in immunosuppressive treatment, in a patient whose disease activity elsewhere has improved, pulls in the opposite direction from an attribution to the disease. The distinction changes treatment completely, since one path escalates immunosuppression and the other reduces it.

Primary psychiatric illness stays on the list throughout, supported by a psychiatric history predating the lupus, a family history, and a typical course. So does the ordinary burden of a relapsing systemic illness with a visible skin component and a treatment with visible adverse effects: demoralisation and adjustment disorder in that setting are a person responding to their circumstances rather than the disease reaching the brain, and are treated as such.

Finally, lupus is one member of a family of systemic diseases that reach the nervous system, and the discriminating features are what separate them at the bedside. The plate accompanying this section sets lupus beside the other systemic autoimmune presentations this chapter teaches, neurosarcoidosis, central nervous system vasculitis, the antiphospholipid syndrome, coeliac disease and inflammatory bowel disease, each with its discriminating feature. Each is taught in full by the sibling section that owns it and is named here only to place lupus in its family.

Six systemic autoimmune conditions on one comparator, because the task is telling them apart
One row each. The middle column is the only one that separates them from each other.

CONDITION	THE DISCRIMINATOR, in one feature	WHAT BRINGS IT TO PSYCHIATRIC NOTICE
Neuropsychiatric lupus	multisystem: skin, joints, kidney and blood counts, all in the one patient	psychosis or a mood episode arising inside an already declared illness
Neurosarcoidosis	cranial neuropathy, with basal meningeal or pituitary involvement	apathy and personality change, with hypersomnia and endocrine mood shift
Primary CNS vasculitis	confined to the CNS, with headache and stepwise deficits on a bland screen	no psychiatric presentation is claimed here: a declared gap, see the text
Antiphospholipid syndrome	thrombosis and pregnancy loss are the carrying features, often recurrent	cognitive complaints in a young adult with a thrombotic or obstetric history
Coeliac disease	gluten-driven enteropathy, and the deficiency load does much of the work	depression and apathy that persist after the mood treatment is optimised
Inflammatory bowel disease	relapsing bowel disease, and the burden tracks disease activity	depression and anxiety that flare with the bowel, plus steroid-driven change

READ DOWN THE MIDDLE COLUMN: that is the only column that closes the differential.
Low mood and cognitive change run across these rows, so that column raises the question but cannot answer it.

● ROUTED OUT: the serological and imaging workup for each of these six, and their immunosuppressive treatment, are not taught on this plate. What is taught here is telling the six apart.

Fig 40.9 Six systemic autoimmune conditions laid on a single comparator so that they can be told apart rather than revised one at a time. The plate exists because the right-hand column is nearly identical down all six rows, and that is the teaching point: depression, anxiety and cognitive change are the common currency of systemic inflammation, so the presentation that brings the patient to psychiatric attention almost never identifies which condition it is. The middle column carries the whole differential, and each entry is chosen to be true of that row and false of the other five. One right-hand cell is deliberately empty of a clinical picture: no source explicitly linking primary central nervous system vasculitis to a psychiatric manifestation was located for this chapter, so the gap is printed rather than filled from a neighbouring condition, and the accompanying text says what to do with it. Note the two structural families the table exposes, that the first four reach the brain directly through inflammation, vasculopathy or thrombosis, while the last two reach it indirectly through nutrition, disease activity and treatment exposure. From memory you should be able to give the single discriminating feature for each of the six, and say why the psychiatric column can raise the question but never settle it.

15.5 Investigation, and what each result can and cannot settle

The investigation logic follows from the absence of a gold standard: every test is chosen for what it excludes or supports, and none to confirm. Serology and disease activity markers establish that lupus is present and give some sense of activity elsewhere, which feeds the temporal axis and nothing more. Structural imaging looks for a focal lesion, for ischaemic change and for an alternative structural explanation, and a normal scan is entirely compatible with a severe diffuse syndrome, which is the most misread result in this topic.

Lumbar puncture here is done principally to exclude infection, and its findings in the disease itself are non-specific. Electroencephalography supports or refutes a seizure explanation for an episodic disturbance and characterises an encephalopathy without attributing it. Neuropsychological assessment characterises a cognitive deficit, which is valuable for

rehabilitation and for tracking change, and does not attribute it. Where autoimmune encephalitis is genuinely in question, that workup belongs to the neurology and psychiatry interface notes.

This chapter holds no row carrying an antibody to syndrome association, an imaging finding to syndrome association, or the performance of any investigation in this population, so none is asserted. That is a real gap rather than a stylistic omission, and the antibody associations should be read from a primary source and held separately from what is taught here.

RULE **No investigation closes the question, because no diagnostic gold standard exists for these syndromes.** Order tests to exclude the competing explanations and to characterise the syndrome, and expect to make the attribution on the balance of what you have. An answer promising a confirmatory test is wrong on the central fact of the topic.

15.6 Treatment, and the recommendation this section carries

Once a syndrome has been attributed to the disease, treatment has three limbs, chosen by mechanism rather than by phenomenology. The **inflammatory limb** is immunosuppression, applied when the syndrome is judged to arise from active disease. The **thrombotic limb** is antithrombotic treatment, applied when the mechanism is judged to be vascular occlusion; the antiphospholipid syndrome that most often drives it is owned by the sibling section, so this section names the limb and stops. The **symptomatic limb** is psychiatric and neurological treatment of the syndrome itself, running alongside either of the others rather than after them.

The guideline recommendation this section carries is precise about its population and about its comparison, and both matter more than the drug names. In adult patients with lupus-related severe, acute neuropsychiatric manifestations, the guidance is to **use GCs plus CYC over GCs alone or GCs plus RTX**. The abbreviations are the guideline's own and stand for glucocorticoids, cyclophosphamide and rituximab.

Read the population before using the recommendation, because it carries restrictions a candidate quoting it loosely will drop. It concerns adults, so it says nothing about children. It concerns severe manifestations, so it says nothing about the many mild ones this container also holds. It concerns acute manifestations, so it says nothing about chronic or slowly progressive syndromes, and a chronic cognitive syndrome is precisely what psychiatry is most likely to be asked about. Quoted as though it applied to everyone with this label, it becomes an error of scope.

What this section does not carry is equally worth stating: no dose, no route, no duration, no schedule, no maintenance strategy, no comparison for the mild or chronic presentation, and no statement about what to do when the first choice fails. Organ-adjusted prescribing for renal involvement belongs to the psychopharmacology notes, and remembered regimens in this topic diverge sharply between sources.

15.7 The psychiatric limb, and the boundaries around it

Psychiatric treatment of an attributed syndrome runs alongside the immunosuppression and is not deferred until it works. An acute confusional state is managed as an acute confusional state, by the pathway taught in the delirium notes, while the cause is treated. A psychotic or affective

syndrome is treated symptomatically with two standing cautions in view: renal and haematological involvement constrain the choice of agent, quantified in the psychopharmacology notes rather than here, and the possibility that treatment of the disease is contributing keeps the drug review open.

Two further boundaries belong in an answer. Capacity is assessed as it is anywhere else, and the only things this section says about it are that it is decision-specific and time-specific, which matter unusually often in a fluctuating encephalopathy; the doctrine and the statutory test belong to the mental health legislation notes. Adherence deserves explicit attention, because the psychiatric syndrome sits on top of a demanding medical regimen, and a cognitive syndrome never named will be recorded as non-compliance.

The rehabilitative and psychological work is most often omitted and has the longest reach. A cognitive deficit that has been characterised can be worked with, at home and at work, whether or not it is reversible, and the demoralisation that follows a burdensome illness responds to being named and addressed rather than to being reclassified as a manifestation of the disease.

15.8 Answering the question

The question is almost always some form of "a patient with known lupus presents with a psychiatric syndrome, discuss". Each row below is a decision rather than a heading, so an answer that walks the rows is reasoning aloud rather than reciting.

DECISION	WHAT MOVES IT	WHAT CANNOT SETTLE IT	WHERE THE DETAIL IS TAUGHT
Which syndrome is this	Phenomenology, course, examination, cognitive profile	Serology, which speaks to the disease and not the syndrome	Here
Is it attributable to the disease	Timing against disease activity, absence of a competing explanation, atypical features	Any single investigation, since no confirmatory test exists	Here
Is it infection, metabolic or hypertensive	Cultures, biochemistry, blood pressure, imaging, lumbar puncture	The presence of lupus, which raises infection risk	Here, and the delirium notes
Is it the treatment rather than the disease	Onset after a treatment change, disease quiescent elsewhere, the drug review	Severity, which is uninformative about cause	The routing section owning the corticosteroid construct
Inflammatory or thrombotic mechanism	Focal against diffuse presentation, vascular imaging, antibody profile	Response to a treatment not yet given	Here for the limbs, the sibling section for the antiphospholipid syndrome
What is given, and how is it prescribed	Severity and acuity of the manifestation, and the age group	A remembered dose or regimen	Here for the recommendation, psychopharmacology for prescribing

IN INDIA

No Indian row was located for this item, and the shortfall is declared rather than filled: no Indian cohort, no Indian frequency estimate, no Indian figure for the cost or availability of either named agent, and no Indian guideline statement on this manifestation. The recommendation this section does carry comes from a Latin American clinical practice guideline, so its transferability to an Indian public hospital is a judgement rather than a given, and a candidate who names that is on stronger ground than one who quotes it as though it were local. What is safe to say in an Indian answer is structural: infection carries the most weight of the exclusions where the background burden of chronic infection is high, immunosuppression is decided jointly with rheumatology, and the rehabilitative limb is the one most likely to be dropped where follow-up is difficult.

PITFALL

Answering as though attribution were a formality on the way to treatment. An answer spent on immunosuppression has skipped the contested part, and also the part that protects the patient, since the competing explanations are commoner and more immediately reversible than the manifestation being reached for. Spend the answer where the difficulty is.

GOOD TO REMEMBER

Called to a ward for a patient with lupus and a new psychiatric presentation, give the medical team your view on whether the disease has reached the brain as a judgement rather than a finding, and record what would change it. An attribution offered in that form can be revised without anybody losing face; one asserted as a finding is very hard to withdraw.

MNEMONIC

Syndrome, Sepsis, Salts, Script, Then Attribute

Characterise the syndrome, exclude sepsis and infection generally, check the salts and the rest of the metabolic screen, read the script for drugs that could be doing it, and attribute to the disease last. Everything before the last step is faster, commoner and more reversible than what comes after it.

No gold standard	Attribution	GCs plus CYC
FOR DIAGNOSING THE NEUROPSYCHIATRIC SYNDROMES OF LUPUS	WHAT DIAGNOSIS ACTUALLY CONSISTS OF HERE	PREFERRED OVER GCS ALONE OR GCS PLUS RTX, ADULTS, SEVERE ACUTE MANIFESTATIONS

There is no test that confirms neuropsychiatric lupus, so diagnosis is an attribution built from timing, exclusion and the character of the syndrome, and the exclusions that matter most are infection, metabolic disturbance and the treatment itself. Treatment then splits by mechanism into an inflammatory limb, a thrombotic limb and a symptomatic limb, and the one recommendation this section carries applies to adults with severe, acute manifestations and to nobody else.

16 · Neurosarcoidosis, CNS vasculitis and antiphospholipid syndrome with psychiatric features

Neurosarcoidosis, primary vasculitis of the central nervous system and antiphospholipid syndrome are filed together as rare systemic causes of psychiatric presentation, and that filing is very nearly the only thing they share. Each arrives in psychiatry by a different route, each is settled by a different kind of evidence, and for one of them the psychiatric literature this text could locate is thin enough that saying so is the correct answer rather than an embarrassment. This section installs what is actually sourced for neurosarcoidosis, what is actually sourced for antiphospholipid syndrome, and a worked method for answering on primary central nervous system vasculitis without inventing a phenotype for it. The rare causes are examinable for exactly this reason: they are where a candidate's discipline about evidence becomes visible inside a single paragraph.

16.1 Three names, three different examination problems

The grouping is administrative. A syllabus has to put the uncommon systemic causes somewhere, and the shared shelf tempts a tired candidate into one undifferentiated answer about inflammation reaching the brain. That answer earns very little. The examiner assumes you accept that systemic autoimmune disease can affect the mind, and is asking something else: whether you can tell these conditions apart with a patient in front of you, and whether you know which parts of the answer are yours to give.

It helps to sort them on a single axis: where the discriminating information sits. For **neurosarcoidosis**, the sourced material puts discriminating information inside the mental state itself, because the described presenting picture assembles features that a purely functional psychosis would not usually assemble in the same person at the same time. For **antiphospholipid syndrome**, the sourced material does the opposite: the reported psychiatric syndromes are the ordinary syndromes of general adult psychiatry, so the mental state is precisely where the discrimination is not. For **primary central nervous system vasculitis**, the discriminating information sits in a specialty literature that this text could not quote, and that changes the shape of an honest answer without making the condition irrelevant.

Those are different problems and they take different answers. The first is a recognition problem: can you notice, in a mental state you are already examining, a combination that sends you outside psychiatry. The second is an attribution problem: with a systemic diagnosis on the file and an ordinary psychiatric syndrome in front of you, how much may you claim about the relation between them. The third is an evidence problem: what do you say when the psychiatric claim you would like to make is one your sources will not carry. A candidate who answers all three with one paragraph about immune-mediated brain injury has told the examiner that the differences are invisible to them. The generic immune-to-brain framework belongs elsewhere in this volume and is deliberately not rebuilt here.

16.2 Neurosarcoidosis: a picture that can arrive before the diagnosis

The sourced statement for neurosarcoidosis is a statement about presentation. Reported neuropsychiatric manifestations include **confusion, dysfunction in orientation or memory, hallucinations, delusions and catatonia**, and these are described as manifestations that can be the presenting features of the condition. Read the load-bearing word carefully. Presenting means the psychiatric picture can be the first thing anyone sees, which in turn means the referral may reach a psychiatrist while the systemic diagnosis is still unmade, unsuspected and absent from the file you are handed.

That is the examinable point, and it is about sequence rather than rarity. Were the systemic diagnosis always established first, this condition would be a footnote in liaison practice, because the medical team would carry it. Because the neuropsychiatric picture can come first, the recognition burden lands on psychiatry, during an ordinary assessment of what looks like a first episode of psychosis.

The discriminating feature is the combination rather than any single item. The sourced list runs across layers that psychiatric nosology usually keeps apart: a confusional and cognitive layer (**confusion, dysfunction in orientation or memory**), a psychotic layer (hallucinations and delusions), and a motor and behavioural layer in the form of catatonia. A patient in whom orientation and memory are disturbed, psychotic phenomena are present, and catatonic signs are observable has a picture crossing those layers at once. Crossing them is the signal, and the response is a decision to widen the question beyond the psychiatry room rather than a diagnosis made inside it.

Be equally precise about what the source does not license. The material behind that list is a report of a single presentation, and its framing describes these manifestations as rare. A statement of that kind establishes possibility. It does not establish frequency, it does not rank the features against one another, and it supplies no proportion of patients in whom any of them occurs. So there is no percentage in this section, and a candidate who supplies one has supplied it from memory. The marks here are for recognising the pattern and routing the patient, never for a prevalence figure.

RULE **A presenting feature is not a common feature.** The sourced claim about neurosarcoidosis licenses recognition: you keep the condition available when a picture crosses cognitive, psychotic and catatonic layers together. It licenses no statement about how often that happens, and none that this presentation is typical of the disease. State the possibility, then say what you would do about it.

16.3 What psychiatry contributes when neurosarcoidosis is the question

Once the question is open, the psychiatrist's contribution is a description good enough for someone else to act on. Record phenomenology in the patient's own words rather than in summary labels, because the summary label is what a later reader will doubt. Record catatonic phenomena as observed signs with the manoeuvres that elicited them, so that a neurologist reading your entry can see what you saw. Date the cognitive change against the psychiatric history: a disturbance of orientation and memory that arrived with the psychotic phenomena

reads very differently from one that has been present for years, and only the psychiatric assessment is positioned to establish that ordering.

The rest of the pathway belongs to other owners, and naming them accurately is part of a complete answer. Imaging and cerebrospinal fluid evaluation of a suspected inflammatory brain process, with immunotherapy decisions, sits with the neurology and psychiatry interface material. Assessment and management of delirium as a syndrome, including its bedside instruments, sits with the delirium and focal brain syndromes material. Corticosteroid-related psychiatric disturbance, relevant once systemic treatment begins, is owned by its own routing section and by the endocrine and metabolic material. Adjustment of psychotropic prescribing for organ impairment sits with the psychopharmacology material. Capacity, if it arises, is decision-specific and time-specific, and the statutory test belongs to the mental health legislation material.

Routing is not evasion. An answer that names the right owner for each limb of the work shows the examiner how you manage a shared patient, and in particular that you know how each team assumes the other has already done the thinking.

TRAP **Recording catatonia as a diagnosis rather than as a set of signs.** Once catatonia is entered as a conclusion, the file behaves as though the question has been answered, and the systemic question quietly stops being asked. Candidates fall into this because catatonia feels like an endpoint: it is a named syndrome, it has a management pathway, and naming it produces the relief of having decided something. Written as observed signs instead, it stays what it is, which is a finding that still needs an explanation.

16.4 Primary central nervous system vasculitis: a declared gap, and how to answer it

Here this text has to say something uncomfortable and say it precisely. What was sought was a self-contained quotation from a primary study, a guideline or a systematic review that explicitly links primary central nervous system vasculitis to a psychiatric manifestation. No such quotation was located for this fragment. That is the gap, stated with its exact boundaries, and the boundaries matter: the search was for an explicit link to a psychiatric manifestation, quotable on its own, in a source of that standing.

Note what the gap is not. It is not a finding that the condition has no psychiatric presentation, and converting it into that claim would be a worse error than the one it replaces, because a categorical negative asserts far more than a failed search can support. Absence of a located quotation is a fact about this text's sources. It is not a fact about the disease.

Say also whose claim it would be. A statement linking this condition to a specific psychiatric syndrome would rest on the authors of peer-reviewed clinical studies of the condition and on the neurology and rheumatology specialty literature, where cohorts of these patients are described. Psychiatry cannot assume it by analogy with the conditions on either side of it in this section. The specialty that follows these patients owns the phenotype, and the honest position is to defer to it by name rather than fill the space.

Clinically the gap changes little about what you do and a great deal about what you write. The condition stays available as an inflammatory cerebrovascular possibility in a patient whose

presentation has crossed into neurological territory, and its evaluation sits with neurology as the neurosarcoidosis workup does. What changes is that you refrain from attaching a named psychiatric syndrome to it in a case note, a referral letter or an examination answer, because you would be attributing on authority you lack. You describe the presentation, record why the possibility is open, and let the specialty that can settle it settle it.

This is worth practising as a sentence, because candidates who have never said it aloud tend to bluff instead. A usable form runs: the psychiatric phenotype of this condition is not something I can state from a specific source, the neurology and rheumatology literature holds it, and my contribution is the described presentation and the referral. That answer is short, it is true, and it displays the judgement that separates a safe registrar from an unsafe one. An invented list of psychiatric features, delivered fluently, is the failure mode this discipline exists to prevent.

TRAP **Filling the vasculitis gap with borrowed lupus material.** Neuropsychiatric systemic lupus erythematosus has a rich catalogue of manifestations, it is taught in its own section of this chapter, and it sits close enough in a candidate's mind that its phenotype slides across to fill an empty space. It is a borrowed answer about a different disease, and an examiner who knows the vasculitis literature is thin will hear the borrowing at once. Name the lupus material as the neighbouring topic it is, and leave its content where it belongs.

16.5 Antiphospholipid syndrome: a psychiatric range that reads like general psychiatry

The sourced position for antiphospholipid syndrome is that a range of psychiatric disorders has been reported in these patients, and the reported range includes **psychosis, mania, depression, bipolar disorders and schizophrenia**. Read that list as a whole before reading any item in it. Taken together, **psychosis, mania, depression, bipolar disorders and schizophrenia** are close to a roster of general adult psychiatry. That is the discriminating feature of this condition, and it is a negative one in a useful sense: there is no distinctive mental state signature here to look for.

The consequence follows directly. If the reported psychiatric range is indistinguishable by name from primary psychiatric illness, the mental state examination cannot be the instrument that raises the possibility, and the discrimination has to come from elsewhere in the record. The syndrome is named for an antibody, which tells you that serology and the general medical history are where the question is settled, and that the settling belongs to the rheumatology and haematology literature. This text carries no sourced criteria for the condition, so it prints none, and a candidate who prints a criteria set from memory is doing the thing this method exists to stop.

Be careful with the verb as well. The sourced claim is that these disorders have been reported in patients with the condition. Reported in is an observation of co-occurrence in a described population, rather than a causal claim about the antibody producing the syndrome, and it carries no proportion. So a psychiatrist assessing a patient who already holds this systemic diagnosis has reason to keep the association in view while treating aetiology as unsettled. The material behind the claim is a narrative review drawing on a long publication window: a reasonable authority for the statement that a range has been reported, and a poor one for anything quantitative. There is no figure in this section either.

What this leaves you with is a clinical posture rather than an algorithm. In a patient with the systemic diagnosis on file who develops a new psychiatric syndrome, each of the easy errors is a real one. Treating the psychiatric presentation as unrelated by default discards information the record already holds. Attributing it to the systemic condition by default converts a reported association into a diagnosis and ends an assessment that had barely begun. The defensible position lies between them: keep the systemic diagnosis as a live question in the formulation, put that question in writing, and manage the psychiatric presentation on its own terms while it stays open.

RULE **Reported in is not caused by.** This distinction carries more marks across the chapter than any single disease fact, because almost every systemic association in liaison psychiatry is reported at the level of co-occurrence. Say what your source says, with the verb your source uses, then say separately what you would do clinically. An answer that keeps the moves apart reads as trained. An answer that fuses them reads as memorised.

16.6 Telling them apart in one pass

The plate for this section places the disease-specific presentations beside one another, including the neighbouring topics this section deliberately leaves to their owners: neuropsychiatric lupus, and coeliac disease and inflammatory bowel disease with the gut-brain axis. Use it as a recall grid in revision, working across the row for each condition and asking what makes that condition itself rather than the one above it. The table below does something complementary: it records where each condition in this section stands in evidence, which governs how firmly you may speak.

CONDITION	WHAT THIS TEXT CAN STATE	WHERE THE REMAINDER IS HELD
Neurosarcoidosis	That **confusion, dysfunction in orientation or memory, hallucinations, delusions and catatonia** can be presenting features, so psychiatry may see the patient first	Frequency, imaging and cerebrospinal fluid evaluation, and immunotherapy: neurology and the interface material
Primary central nervous system vasculitis	A declared gap: no self-contained quotation explicitly linking it to a psychiatric manifestation was located for this fragment	The psychiatric phenotype, if described: authors of peer-reviewed clinical studies and the neurology and rheumatology literature
Antiphospholipid syndrome	That a range of psychiatric disorders has been reported, including **psychosis, mania, depression, bipolar disorders and schizophrenia**	Diagnostic criteria, serology and proportions: the rheumatology and haematology literature

16.7 Answering on the rare causes without over-claiming

A complete answer here has a recognisable shape, worth rehearsing because the content is thin enough that structure carries the marks. Five moves carry it.

- Name the three conditions separately, and refuse the single undifferentiated paragraph about inflammation reaching the brain that the shared shelf invites.
- Neurosarcoidosis: give the discriminating feature as a presenting picture crossing cognitive, psychotic and catatonic layers together, and say that the response is to widen the question rather than to make the diagnosis in the psychiatry room.
- Antiphospholipid syndrome: give the discriminating feature as a reported psychiatric range that reads like general adult psychiatry, so the discrimination lives outside the mental state, in the record and the serology.
- Primary central nervous system vasculitis: give the discriminating feature as a psychiatric literature this text could not quote, held by the specialties that follow these patients, and name it as theirs rather than filling the space.
- Mark explicitly which parts of your answer rest on a source and which rest on clinical reasoning, then route the workup to its owner.

Recognition	**Attribution**	**Evidence**	**Reported, not caused**
THE NEUROSARCOIDOSIS PROBLEM	THE ANTIPHOSPHOLIPID PROBLEM	THE VASCULITIS PROBLEM	THE VERB THAT CARRIES THE MARKS

The first three tiles are the three problems this section separates, one to each condition, and a candidate who can say which of the three a question is asking has already done most of the discriminating. The fourth is the verb discipline that runs underneath all three.

The shape survives contact with a question you were not expecting, which is its real value. The neighbouring topics in this chapter, paraneoplastic neuropsychiatric syndromes, neuropsychiatric lupus, immunosuppressant neurotoxicity after transplantation, and the gut-brain topics, are owned elsewhere in this volume and each carries its own sourced material. Name them as neighbours when a question invites the comparison, and resist teaching them inside an answer about these conditions. Examiners read borrowed content as padding, and padding is expensive in a viva.

IN INDIA

No Indian series or Indian guidance on the psychiatric presentation of these conditions was located for this fragment, and that is a limit of this text rather than a claim about Indian practice or the Indian literature. What travels regardless of setting is the referral, and it is worth writing well. Specialist assessment may be separated in time from the presentation you witnessed, so the phenomenology, the catatonic signs with their eliciting manoeuvres, and the dating of cognitive change against the psychiatric history are things only your entry preserves. Where rheumatology, neurology or immunology assessment means travel and delay for the family, the referral letter is what keeps the question alive until they are seen.

GOOD TO REMEMBER

Ask the systemic question at the moment a psychiatric picture stops behaving like one, rather than at the end when everything else has failed. For neurosarcoidosis the trigger is a presentation crossing cognitive, psychotic and catatonic layers together. For antiphospholipid syndrome it comes from the record rather than the mental state, because the reported syndromes are the ordinary ones. For primary central nervous system vasculitis it is the neurological unease that would prompt any inflammatory cerebrovascular referral, and your written answer stops short of naming a psychiatric syndrome for it. Describe fully, attribute cautiously, route to the specialty that can settle it.

Three conditions, three different problems: neurosarcoidosis is a recognition problem, antiphospholipid syndrome an attribution problem, primary central nervous system vasculitis an evidence problem, so describe fully, attribute cautiously, and route each to the specialty that can settle it.

17 · Coeliac disease, inflammatory bowel disease and the gut-brain axis in psychiatry

Two bowel diseases reach the psychiatry examination for opposite reasons: coeliac disease for what it is associated with, and inflammatory bowel disease for how it behaves over time. Both are usually taught under the single heading of the gut-brain axis, and that heading does them a disservice, because it places an epidemiological observation and a mechanistic theory in the same sentence and lends them the same confidence. This section separates them: what the comorbidity evidence in coeliac disease states and what its design will bear, what the bidirectionality evidence in inflammatory bowel disease states and what it will not bear, what each finding changes in front of a patient, and where the mechanistic account stops being evidence and becomes a story that is merely satisfying. The examinable skill here is discrimination, not enthusiasm.

Several neighbouring topics are named here and left to the sections that own them. Neuropsychiatric systemic lupus erythematosus, neurosarcoidosis, central nervous system vasculitis and antiphospholipid syndrome each have their own home in this chapter, and their manifestations are not repeated here. The general immune-brain vocabulary of inflammatory signalling and sickness behaviour belongs to the earlier notes that own it and is used here only as a name. Steroid-induced psychiatric disturbance is owned by the routing section of this chapter and appears below only as a confound to be excluded before anything else is concluded. Delirium assessment belongs to the notes on stroke, traumatic brain injury and focal brain syndromes. Organ-adjusted prescribing, including any adjustment for malabsorption, belongs to the psychopharmacology notes, so no dose, no interaction ranking and no drug choice is stated in this section.

17.1 Three different claims wearing one name

The phrase **gut-brain axis** is used to make at least three claims that differ sharply in how well they are supported, and most of the confusion in this topic comes from sliding between them

inside a single answer.

- A **co-occurrence claim**: people carrying the bowel diagnosis carry more psychiatric morbidity than some comparison group. This is a question about counts in two populations, and it is the easiest of the three to test.
- A **covariation claim**: within a single person, the state of the bowel and the state of the mind move together over time, and each may carry information about the other's future. This needs longitudinal data and is harder to establish.
- A **mechanistic claim**: a specified biological pathway carries the influence in one direction or both. This is the claim the informal literature makes most often and the one it supports least well.

The two diseases in this section illustrate the first two claims respectively. Neither of the findings taught below is a mechanistic claim, and keeping that straight is most of the work an examiner is testing.

17.2 Coeliac disease: the comorbidity signal, and the shape of it

The finding to carry into the examination comes from a systematic review that compared people with the diagnosis against healthy controls. It **identified a significant increased risk for autism spectrum disorder, attention deficit hyperactivity disorder, depression, anxiety and eating disorders in patients with coeliac disease compared with healthy controls**. That sentence rewards slow reading, because each of its parts constrains what may afterwards be done with it.

The comparator is **healthy controls**, and that one phrase governs the whole interpretation. A comparison against healthy people cannot separate what belongs to coeliac disease specifically from what accompanies having any chronic diagnosis at all: a restrictive treatment, a lifetime of appointments, a label, and the ordinary demoralisation of managing a condition that never resolves. A candidate who says the finding demonstrates a disease-specific psychiatric vulnerability has reported a comparison that was never made. The defensible formulation is that psychiatric morbidity is raised relative to people without the disease, and that the design does not tell us whether it is raised relative to people carrying a different chronic illness of similar burden.

The outcome list is broad rather than narrow. It **spans neurodevelopmental, affective, anxiety and eating-disorder categories** rather than picking out a single syndrome, and breadth of that kind fits a general elevation of psychiatric morbidity better than it fits a specific lesion producing a specific disorder. That does not make the finding uninteresting. It makes it a reason to take a full psychiatric history in this population rather than a reason to go looking for one particular condition.

The design is a synthesis of observational comparisons. It establishes that the categories co-occur more often than expected, and it establishes nothing at all about which of them came first. The

rival explanations that survive a design of that kind are set out below, and reciting them is the cheapest mark available in this topic.

TRAP **Converting an association list into a causal list.** The review reports raised risk across several diagnostic categories against healthy controls. Candidates routinely restate this as the bowel disease causing the psychiatric disorders, then reason forward to a treatment implication that no observational synthesis can support. The finding tells you whom to assess carefully. It does not tell you what to attribute, and it does not tell you what to treat first.

17.3 What the coeliac finding actually changes in the room

Two situations arise, and they call for different work. In the first, the bowel diagnosis is already made and the person is referred with a psychiatric presentation. Here the finding does one useful thing: it removes any temptation to treat the psychiatric picture as an obvious reaction to illness that needs no further characterisation. Because the associated categories **include neurodevelopmental as well as affective and anxiety conditions**, a history that stops at mood in the last few months is too short. The developmental history, the school history and the pattern of attention and social functioning across the lifespan belong in the assessment, and they are the parts most often skipped when a medical diagnosis is already on the front of the file.

The entry in that list which most changes a consultation is **eating disorders**, because the management of coeliac disease is itself dietary exclusion. A medically required restriction and a symptomatic restriction occupy the same behavioural territory, and each can conceal the other. A person whose intake has narrowed far beyond what the bowel disease requires may be presented by the family, and even by the treating team, as unusually adherent. The clinically useful question is not whether the person avoids particular foods, which the treatment already dictates, but what the avoidance is for, what happens when it is breached, whether weight and shape have entered the reasoning, and whether the restriction has spread past the medical boundary into foods it never covered. That distinction is doable in the consultation and is where the psychiatric contribution is greatest.

In the second situation the bowel diagnosis has not been made and the temptation is to reason towards it from the psychiatric presentation. Here the section declares a limit rather than manufacturing guidance. No evidence row available to it grounds a screening rule, a case-finding threshold, an at-risk group worth testing or any performance figure for a diagnostic test, so none is printed. What survives without a row is procedural, and it is enough for most examinations: gastrointestinal symptoms, unexplained weight loss and unexplained nutritional deficiency belong to the referring physician, and the right move is a clear referral question rather than a test ordered from the psychiatry clinic.

17.4 Inflammatory bowel disease: two arrows, and the modal verb that carries them

For inflammatory bowel disease the claim to hold is different in kind. A systematic review and meta-analysis concluded that **bidirectional effects of the brain-gut axis are present in inflammatory bowel disease and may influence both the natural history of the disease and psychological health**. This is a covariation claim, not a co-occurrence claim, and it is the more

clinically interesting of the two because it concerns what happens next rather than what is present now.

Unpacked, **bidirectionality** is two arrows, and the conclusion adds a qualifier. The first arrow runs from the bowel to the mind: the state of the disease bears on psychological health. The second runs from the mind to the bowel: psychological state bears on the natural history of the disease. A complete answer names both, because naming only the first reduces psychiatric symptoms to an understandable reaction and leaves the psychiatrist with nothing to offer except sympathy, while naming only the second overstates what any single service can influence.

The qualifier matters as much as the arrows. The conclusion is that these effects **may influence** disease course and psychological health. It is a statement that bidirectional effects are present and that influence is possible. It is not a magnitude, it is not a treatment claim, and it is emphatically not a finding that treating a mood disorder alters the course of the bowel disease. That last step is the one candidates take without noticing, usually in the final sentence of an otherwise careful answer, and it converts a defensible position into an unsupported one.

TRAP **Promoting a bidirectional association into a therapeutic promise.** "The brain-gut axis is bidirectional, therefore treating depression will improve the bowel disease" reads as a natural inference and is a different claim from the one the evidence makes. Treat the psychiatric disorder because it is a disorder and it causes suffering. Do not sell it to the patient or to the gastroenterology team as disease-modifying therapy.

17.5 Working both arrows at the bedside

Bidirectionality becomes concrete as soon as it is turned into a sequence of consultation moves. Because **the effects run in both directions**, the first task is to work out where in the cycle the person currently sits. The psychiatric history on its own will leave that question open.

Begin with disease activity. Before attributing fatigue, poor concentration, disturbed sleep and reduced appetite to a mood disorder, establish with the treating team whether the bowel disease is currently active, when it was last assessed, and what the plan is. The activity measures used by the gastroenterology service are theirs to interpret and are not named or scored here; what the psychiatrist needs is the team's answer, in writing, to the question of whether this is active disease. A psychiatric formulation written without that answer is a guess about which arrow is operating.

Next exclude the treatment as the explanation. Corticosteroid-induced psychiatric disturbance is taught in full by the section of this chapter that owns it, and is named here only so that it is excluded before anything is attributed to the axis. Establish what the person is taking, what was started or stopped in the weeks before the change in mental state, and whether the timing fits. A mental state change that tracks a treatment change is a treatment problem until proven otherwise, and no amount of gut-brain reasoning displaces that.

Then characterise the illness experience rather than assuming it. Urgency, incontinence, nocturnal symptoms and the resulting narrowing of where a person is willing to go are facts that patients rarely volunteer, because they are humiliating. Asking about them directly is the part of

the interview that most often turns a generic depressive episode into a specific and modifiable set of avoidances.

RULE **Establish disease activity and treatment timing before you formulate.** In a person with inflammatory bowel disease, low mood, fatigue and poor concentration have at least four candidate origins: active disease, the treatment, the restriction the illness imposes on daily life, and an independent psychiatric disorder. These are not mutually exclusive and they demand different responses. Writing to the gastroenterology team with a specific question, and waiting for the answer, is faster than treating the wrong one.

17.6 Mechanism, and where a careful answer stops

This is the part of the topic where overreach is routine, so the boundary is worth stating plainly. The mechanistic literature on gut-brain communication is large and is written with more confidence than its evidence carries. Several families of explanation are commonly offered: immune and inflammatory signalling, the gut microbial community and its products, nutritional and absorptive consequences of the bowel disease, the effects of treatment, shared genetic or developmental liability acting on both systems, and the ordinary psychological load of a chronic, unpredictable and stigmatising illness. These are named here as categories only.

No evidence row available to this section grounds any specific mechanism in either disease, and none is asserted here. That is a declared gap, not an oversight: the two findings this section owns are an association in coeliac disease and the presence of bidirectional effects in inflammatory bowel disease, and neither identifies a pathway. A candidate asked to explain the axis mechanistically should name what is proposed, mark it as proposed rather than established, and then give what the evidence does support, which is the two claims above.

Four alternatives should be recited whenever an association in this field is offered as though it were a mechanism. The association may run in the reverse direction from the one assumed. A third factor may produce both. Detection may be differential, so that people already in one care system are more thoroughly investigated in the other. And the association may be entirely real yet wholly mediated by the burden of chronic illness, which is a psychological explanation and needs no biology at all. An answer that lists these four has demonstrated the exact competence being examined.

TRAP **Reaching for the microbiome because it is available.** A fluent mechanistic paragraph is easy to produce in this topic and is the commonest way of converting a good answer into a suspect one. Fluency is not evidence. Say what is proposed, mark it as proposed, and spend the remaining time on the two findings you can actually defend and on what they change in the consultation.

17.7 Where these two sit among the systemic presentations

It is worth placing coeliac disease and inflammatory bowel disease alongside the other systemic conditions that reach psychiatry, because they belong to a different column. The autoimmune conditions taught elsewhere in this chapter present as direct neuropsychiatric manifestations, in which the illness process itself produces the psychiatric picture and the diagnostic task is to recognise it. The two bowel diseases here are not taught that way. Their psychiatric relevance is

an elevated burden of common psychiatric disorders in coeliac disease and a bidirectional relationship with disease course in inflammatory bowel disease, which is a question of comorbidity and prognosis rather than of a direct neuropsychiatric syndrome.

The systemic autoimmune comparator plate sets each systemic presentation in this chapter beside its discriminating feature, so the conditions can be told apart in the moment rather than merely listed. Read it as a routing aid: the discriminating feature is what should redirect the assessment when it is present. The two entries relevant here are the ones whose discriminating feature is not a neuropsychiatric sign at all, and noticing that asymmetry is the point of putting them on the same plate.

That plate is printed with the neuropsychiatric lupus section, where it places these two conditions beside the other four, and it is the fastest way to revise the discriminating features together.

17.8 What a complete answer contains

A structured answer on the gut-brain axis contains five things in this order: the definition with the three claims separated; the coeliac finding stated with its comparator, which is the part most answers drop; the inflammatory bowel disease finding stated with both arrows and the qualifier intact; the design limitations, meaning direction, shared liability and differential detection; and the consultation consequences, meaning a lifespan history in coeliac disease, the required-versus-symptomatic restriction question, and disease activity and treatment timing established before any psychiatric attribution in inflammatory bowel disease.

QUESTION	COELIAC DISEASE	INFLAMMATORY BOWEL DISEASE
Type of claim the evidence supports	Co-occurrence against a comparison group	Covariation over time, in both directions
What is stated	**Increased risk of autism spectrum disorder, attention deficit hyperactivity disorder, depression, anxiety and eating disorders against healthy controls**	**Bidirectional brain-gut effects are present and may influence disease course and psychological health**
What is not stated	Direction of causation, disease specificity against another chronic illness, any mechanism	Any magnitude of effect, and any claim that psychiatric treatment alters the natural history
First move in the consultation	Full lifespan psychiatric history, developmental included, and the required-versus-symptomatic restriction question	Establish current disease activity and treatment timing with the treating team, in writing
Commonest error	Reading the association list as a causal list	Promising disease modification from psychiatric treatment

IN INDIA

No Indian prevalence figure, service audit or cost estimate for either disease appears in this section's evidence rows, so none is printed; the shortfall is declared rather than filled. What can be said is organisational. The gastroenterology service is usually in a different building from psychiatry and often in a different institution, so the liaison here is a written one, and it should carry a specific question rather than a request for an opinion: is the disease active, what changed in the treatment and when, and what is the plan for the coming months. The dietary restriction that treats coeliac disease is negotiated inside a shared kitchen, with food prepared for the household at once, so an adherence discussion aimed only at the patient misses where the decisions are actually made. And where a person travels a long distance to a tertiary clinic, align the psychiatric and gastroenterology reviews to the same visit wherever the services can manage it, because a review that needs a second journey is a review that will not happen.

GOOD TO REMEMBER

In coeliac disease, widen the history: the associated categories reach into development and into eating behaviour, so a history that covers only the last few months of mood is the wrong shape. In inflammatory bowel disease, narrow the question: find out whether the disease is active and what the treatment has been doing before you attribute anything to the axis. In both, treat the psychiatric disorder because it causes suffering and is treatable, and resist the pull to sell it to the patient or to the referring team as a way of changing the bowel disease.

MNEMONIC

Widen, then narrow

Widen for coeliac disease, because the associated conditions cross the lifespan and the eating-behaviour boundary. Narrow for inflammatory bowel disease, because the useful question is which arrow is operating right now, and that is answered by disease activity and treatment timing rather than by mechanism.

Three, not one	Comorbidity, not cause	Possible, not measured	Proposed, not established
CLAIMS INSIDE THE ONE PHRASE	WHAT THE COELIAC REVIEW SUPPORTS	WHAT THE QUALIFIER LEAVES OPEN	HOW MECHANISM MUST BE MARKED

Coeliac disease gives you a comorbidity fact and no direction; inflammatory bowel disease gives you a direction in both senses and no magnitude. Everything beyond those two sentences that sounds mechanistic is a proposal, and saying so out loud is what separates a careful answer from a fluent one.

18 · Burn injury and its psychiatric course

Burn injury is the one medical event in this chapter where the psychiatric question runs forwards and backwards: what the injury has done to the patient, and what the patient may have been doing when it happened. Everything examinable here follows from holding those

apart. This section installs a timeline, because each phase of burn care asks a different question, then a question that sits outside the timeline because it is asked at presentation and concerns the moment before the flame. It also states what this text can source and what it cannot, since burns psychiatry is taught nowhere else in these notes.

18.1 The acute phase: the ward you are called to

The referral comes from a surgical or burns team, and psychiatry walks into an environment organised around fluid balance, wound and airway. The mental state you are asked about is produced by several processes at once, and the acute-phase skill is separating them rather than settling on a label. Distress belonging to the event, distress belonging to the pain, a disturbance of consciousness and the effects of what is given for the pain are all live in the same bed at the same hour, and each responds to something different.

What was sought here, and not obtained in verifiable form, was a quotable statement from a burns or trauma source able to carry it that psychological shock, **dissociation** and **acute stress disorder** may occur in the acute phase after burn injury. State the reason exactly, because it changes what the gap means: of the burns records assembled for this section, several hosts refused automated retrieval, one address no longer resolved, and one quotation could not be found in the document actually served. Each of those is a fact about retrieval, none about burn patients. The claim, if made, would rest on the burns and trauma literature, where acute-phase presentations after burn injury are described from real cohorts, and which psychiatry reads rather than owns.

Without the figure, you describe. Acute distress after a burn injury is observable at the bedside and needs no prevalence attached. Record the phenomenology in the patient's own words, date it against the injury and against each operation, and note whether it fluctuates. That record lets a later clinician settle the question you could not, and it beats a confident figure from nowhere.

One process must be named and handed straight over. Disturbed consciousness, disorientation and fluctuating attention occur in burn patients, and the acute burns bed is one place a psychiatrist meets them. This chapter states that they occur and stops there. Assessment, motoric subtypes and management belong to the Stroke, TBI, delirium and focal brain syndromes notes. The routing is deliberate: a second delirium pathway written inside a burns section is how sibling chapters come to disagree, and a disagreement about an acutely unwell patient is a safety defect rather than an editorial one.

RULE **In the acute phase, ask which process, not which label.** Acute distress, pain, disturbed consciousness and drug effect co-occur in a burns bed and respond to different things. The examinable act is separating them at the bedside and routing each to its owner: the disturbance of consciousness out of this chapter, analgesia to the treating team.

18.2 Pain, and the part of it that is yours

Pain is the largest psychiatric problem of the acute phase and the one most easily handed back whole. Burn pain is unusual in being renewed deliberately: dressing changes, debridement and physiotherapy each produce a further episode of severe pain at a time the patient can predict, so

anticipation becomes part of the pain. That anticipatory limb is psychological and sits inside your remit whether or not the analgesic chart does.

The sourced statement available here is modest and useful. In a systematic integrative review of psychotherapeutic interventions for burns patients, **relaxation breathing techniques are an intervention to be considered for pain and anxiety relief in the care of burns patients**. The pairing is the teaching: the intervention is offered for pain relief and for anxiety relief together, which matches what the ward presents, a patient whose pain and whose fear of it have become one complaint by the third dressing change.

Its practical virtues belong in an answer: it needs no equipment, costs nothing, can be delivered by nursing staff once taught, and can be rehearsed with the relative at the bedside. It is an adjunct, and an adjunct sits alongside analgesia rather than in place of it. Analgesic prescribing stays with the treating team and this chapter prints no drug and no dose. Psychotropic selection and any organ-related adjustment belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes, which this chapter adds no further version of.

TRAP **Handing pain back whole.** Candidates write that analgesia is the surgical team's business, which is true of the prescription and false of the problem. Anticipatory anxiety before a predictable painful procedure is a psychological mechanism with a psychological adjunct, and an answer that omits it surrenders the part of burn pain psychiatry was consulted about.

18.3 Sleep, and where its causes sit

Sleep fails early after burn injury, and it fails for reasons that come from more than one place. What was sought here, and not obtained in verifiable form for the retrieval reasons already given, was a quotable statement that poor sleep in burn patients reflects psychological factors and the hospital environment together. That claim would rest on the burns and trauma literature, and on the ward-based sleep work within it.

The gap costs little at the bedside, because the instruction underneath it survives intact. Take the sleep history so that the sources separate: on one side what the ward is doing, meaning light through the night, noise, observations, position and procedures timed for the day list; on the other what the mind is doing, meaning intrusive re-experiencing, nightmares, worry about appearance and money, and the anticipated dressing change. Ask about each explicitly, because a patient asked only whether they are sleeping badly will say yes and tell you nothing actionable. Then correct the environmental limb first, since a hypnotic prescribed into a bright and noisy bay treats a cause that was never in the patient.

18.4 The subacute phase: appearance and disfigurement

As the patient stabilises, the question changes from survival to what the body now looks like and what that will mean for the life attached to it. Keep these apart. **Disfigurement** is a visible difference in the body, assessable by anyone looking at it. **Appearance concern** is what the person makes of that difference, assessable only by asking them. An assessment that measures the burn and infers the distress has measured the wrong variable, and it will be wrong in either direction.

The sourced finding comes from a multi-centre prospective cohort of adults hospitalised after burn injuries. **Increased psychological flexibility during hospital admission after a burn injury predicts lower appearance concerns over time, although not over the effect of baseline appearance concerns.** Each part carries an examinable point. The exposure was measured during the admission, so the acute admission is a measurement window and not only a treatment window, and the outcome came later, so an acutely measured variable travels forward into a later-phase result. That is why this section is a timeline rather than a list of complications. And the qualification is what an examiner will press: the prediction does not hold over and above baseline appearance concerns, so what the person already thinks about their appearance carries the prognostic weight.

The clinical consequence is direct. Ask about appearance concern during the admission, because it carries the prognostic weight and costs one question, and record it so that later change can be read against something. Be careful, though, describing any intervention aimed at flexibility, because the study supporting the association is the study that qualifies it.

TRAP **Reading a predictor as a treatment target.** A variable that predicts an outcome, then loses its prediction once a baseline is controlled for, is telling you where to look rather than what to change. Candidates keep the first clause of that sentence and drop the second, and the dropped clause is why the finding was reported as it was.

Social re-entry also begins in the ward rather than at the front door, and the work is concrete: rehearsing the first meeting with people who have not yet seen the scars, deciding what the patient will say when a stranger asks, and preparing the family, whose answers to the same question set the tone the patient has to live with.

PHASE	THE QUESTION IT ASKS	WHERE IT IS ANSWERED
Acute	Which process is producing this mental state, and what does each need?	This section, for acute distress and for the psychological limb of pain and sleep.
Acute	Is this a disturbance of consciousness, and how is it assessed and managed?	The Stroke, TBI, delirium and focal brain syndromes notes. Named here, never taught here.
Subacute	What does this person make of the disfigurement, and how closely are they followed?	This section, with its prevalence gaps declared.
Any phase	Which psychotropic, and how is it adjusted for organ impairment?	The Psychopharmacology notes. This chapter adds no further version.
Any phase	Does this patient have capacity to refuse the dressing change?	The Mental health legislation notes. Capacity is decision-specific and time-specific; the test is taught there.
At presentation	How did this burn injury happen, and was it an act?	This section, as an antecedent question.

18.5 Depression and post-traumatic stress: what is missing, and what holds

Here the declared gaps bite hardest. What was sought first was a citable prevalence of post-traumatic stress disorder after burn injury, as a proportion of a defined cohort at stated follow-up points inside the first year. What was sought second was a pooled estimate linking burn extent, as a proportion of **total body surface area** above a stated threshold, to an increased risk of depression expressed as a risk multiple. Neither was obtained in verifiable form, so this text supplies neither the proportions, nor the threshold, nor the multiplier. Each would rest on the burns and trauma literature, on primary follow-up cohorts for the first and a quantitative synthesis of such cohorts for the second.

Now the part that carries marks. A missing figure is not a negative finding, and converting one into the other is a larger error than the one it replaces. To say that post-traumatic stress after burn injury is uncommon, or that burn extent is unrelated to later depression, asserts far more than a failed retrieval can support. The correct sentence is short: the figure exists in the burns literature, I cannot state it from a source, and it does not change what I do with the patient in front of me. It does not, because prevalence tells you how hard to look, never whether this patient has the syndrome. That is settled by assessment, applied during admission, at the transition out of acute care, and across the long rehabilitation tail when the operations have stopped and the attention has thinned.

What holds, sourced and directive, is what to do with a patient who screens badly. European practice guidance for burn care states that **patients with an abnormal psychological profile, including suicidal ideation, should be given special attention, adequately monitored and regularly followed by appropriate mental health professionals**. Notice its shape. It is written around a profile rather than a diagnosis, so it is actionable on the day of the assessment and before any diagnostic conclusion is reached. It names obligations of distinct kinds, attention and monitoring and follow-up, which are commonly collapsed into one vague plan. And it names suicidal ideation inside the profile, which is the hinge into the last part of this section.

TRAP **Turning a missing number into a negative finding.** The candidate who cannot recall a prevalence says the association is weak or the disorder is rare. That is a categorical claim built on a failure of memory, unsafe in exactly the domain where being wrong costs a life. Name the gap, name whose literature holds the figure, and state what you do anyway.

18.6 Rehabilitation, and the return to work

Rehabilitation after a major burn injury is measured in years. Scars mature slowly, contractures develop and are released, and each new operation re-opens the acute and subacute questions in a patient who had begun to leave them behind. The consequence is unglamorous and examinable: psychiatric follow-up must outlast surgical follow-up, and the plan should say so in writing, because attention thins once the wounds are closed and it is then that appearance concern and social re-entry become the whole of the patient's day.

A further gap is declared here. What was sought was a comparison of work-reintegration outcomes between burn survivors whose injury happened at work and survivors injured

elsewhere, with the direction of the difference stated. It was not obtained in verifiable form, and it would rest on the burns rehabilitation and occupational health literature, where such cohorts are followed and outcomes defined well enough to compare.

The instruction underneath it survives the gap: ask where the burn injury happened, because the answer reshapes the rehabilitation problem. When the injury occurred at work, the site of the injury and the site of return are the same place, and that single fact brings in fear of the machine or the flame, the employer's response, blame, compensation and colleagues who witnessed it. A domestic injury raises its own version, since the kitchen is also a place the patient must return to. In Indian practice much work is informal, so returning to work may mean returning to a stove, a field or a roadside stall, with no graded return to arrange and financial pressure that overrides advice framed as a preference.

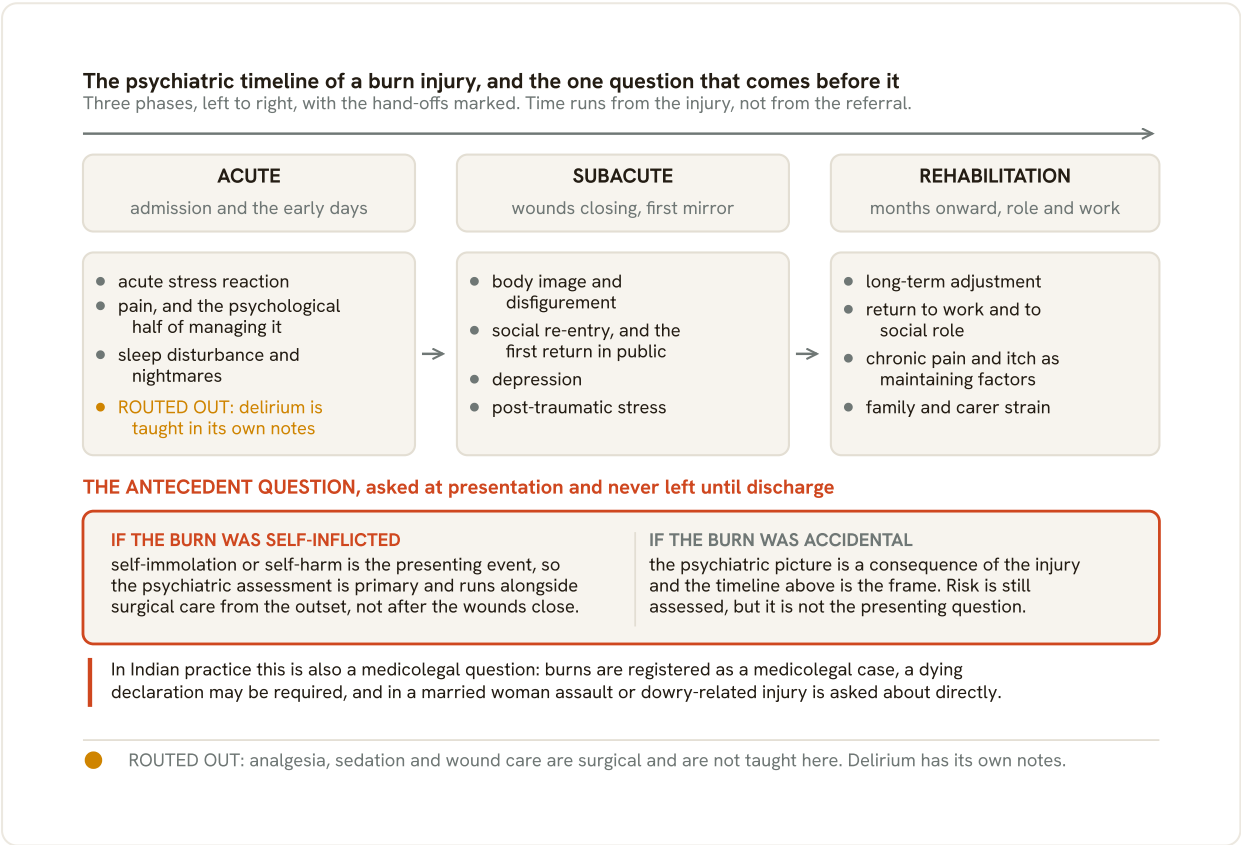


Fig 40.10 The psychiatric course of a burn injury across three phases, with the antecedent question set apart from the timeline because it is not part of it. This is the only plate in the series that teaches burns, so it is built to carry the whole construct rather than illustrate one corner of it. The phases are not merely descriptive: each has a characteristic failure, which is treating acute distress as a disorder before the wounds have closed, missing the subacute period in which body image and the first return in public do most of the psychological damage, and discharging at wound healing as though the rehabilitation phase were somebody else's problem. The block below the timeline is the discrimination that changes management, because a self-inflicted burn makes the psychiatric assessment the primary one from the outset rather than a consequence to be picked up later, and in Indian practice the same question carries a medicolegal weight that an accidental injury does not. From memory you should be able to name the phase-specific problems in order, say where the delirium pathway hands off, and state why self-harm is asked about as an antecedent at presentation rather than as an outcome at discharge.

18.7 The antecedent question, asked at presentation

Everything above treats psychiatric problems as consequences of burn injury. There is a question that runs the other way, asked at presentation rather than at discharge: was the injury itself an

act, and was psychiatric disorder present before the flame rather than after it. That discrimination changes management, and it is the one most often lost when a burns referral is treated as routine liaison work.

It changes management because consequence-work and antecedent-work run on different clocks. Consequence-work is rehabilitative and can be paced. Antecedent-work is risk-work and cannot be: if the burn injury was self-inflicted, the patient is already inside a high-risk period, the method remains available on discharge, the collateral history has to be taken while relatives are still in the hospital, and the discharge destination becomes a psychiatric question. A service that asks how the burn happened only when the wounds are healing has asked its most important question several weeks late.

The Indian relevance is direct, because self-immolation presents to burns units here rather than to psychiatric services, and because the Indian psychiatric literature has examined it. **An Indian report investigated four cases of attempted self immolation, with the psychiatric evaluation done according to DSM-III-R criteria.** Be exact about what that establishes. It is a small clinical series with formal diagnostic assessment applied to it, so it establishes the construct as one Indian psychiatry has studied with criteria rather than impression. Its denominator is small, and using it as though it described a national picture would misuse it.

The same literature shows that self-immolation is heterogeneous rather than a single clinical entity. In an Indian psychosocial study, **22 young people, mostly students, had indulged in this act to protest against the decision of the Government of India**. A public and collective act of protest is a different phenomenon from a private act of self-harm during a domestic crisis, and a formulation built for one will misdescribe the other. So the clinical question is simple to ask and easy to skip: what was the act for. The same physical burn injury can follow any of these.

- A suicidal act, after which the patient is already inside a high-risk period and the method remains available on discharge.
- A public and collective act of protest, which is the phenomenon the Indian psychosocial study above describes.
- An impulsive act inside a family quarrel, a private act during a domestic crisis rather than a public one.
- An accident, where the injury was not an act, though the question of psychiatric disorder before the flame still stands.
- An assault by another person, which this section names as an antecedent and does not follow further.

Those demand different assessments and different collateral sources, which is why the question belongs at presentation rather than in the discharge summary.

Where the assessment finds a self-inflicted injury, ongoing suicidal ideation, or a disturbed psychological picture of any kind, the sourced instruction above applies directly and is worth quoting rather than softening: **special attention, adequate monitoring and regular follow-up**

by appropriate mental health professionals. In a burns unit those obligations have concrete forms. Attention means your assessment happens at presentation, in the surgical bed, written where surgical staff will read it. Monitoring means the nursing observations reflect a psychiatric risk as well as a wound. Follow-up means a named appointment attached to a named person, since a patient who leaves with a dressing clinic date and a vague psychiatric intention will attend the first and lose the second. Where such a patient then refuses dressing changes, capacity becomes the question, and capacity is decision-specific and time-specific; the statutory test is taught in the Mental health legislation and forensic psychiatry notes.

IN INDIA

The burns unit is usually where this psychiatry happens, because the patient occupies a surgical bed and the psychiatric team travels to them. Your assessment competes with a dressing round, your notes are read by surgeons, and the family outside is your collateral source and the carer who will do the dressings at home. Distance and cost shape everything after discharge, so a plan written for a family travelling a long way looks different from one written for a city outpatient. Note what this text could and could not verify. The acute-phase, prevalence and work-reintegration claims were not obtained in verifiable form here, which is a limit of this text rather than a statement about the Indian record. What was verified is Indian, the self-immolation work, and that tells you where Indian psychiatry has looked hardest at burns.

GOOD TO REMEMBER

Ask how the burn injury happened before you ask how the patient is coping with it, because the first question sets the urgency of everything else. In the acute bed, separate the processes and route the disturbance of consciousness to the Stroke, TBI, delirium and focal brain syndromes notes rather than managing it here. Treat anticipatory pain as your work and offer the psychological adjunct alongside the analgesia the team prescribes. Ask about appearance concern during the admission, because it is the early variable that carries the later weight. Where a figure is missing, name the gap and the literature that owns it, then say what you will do anyway.

Antecedent first	**Which process, not which label**	**Asked in the admission**	**A named gap, not a negative**
ASKED AT PRESENTATION, NOT AT DISCHARGE	WHAT THE ACUTE BURNS BED ASKS OF YOU	WHEN APPEARANCE CONCERN CARRIES THE LATER WEIGHT	WHAT A PREVALENCE YOU CANNOT SOURCE BECOMES

Burns psychiatry runs in two directions: ask at presentation whether the injury was an act, then follow the timeline of what the injury has done, separating process from label in the acute bed, treating anticipatory pain as your work, asking about appearance concern during the admission, and naming a missing figure rather than converting it into a negative finding.

Consolidate and apply

19 · Worked vignettes

Every referral in this chapter asks one question in a different accent: is this the medical illness, is this the treatment for it, or is this a psychiatric disorder that would have arrived anyway. The sections before this one supply the content for that triage. What they cannot supply is the experience of running it against a referral somebody else has already answered, where all three possibilities are live at once and the answer is sometimes that the question belongs to another chapter.

Each vignette opens with the referral as it would arrive, works the reasoning in the open, and closes with what an examiner is testing. Where a vignette reaches a specific taught elsewhere, it names the owning section and stops. Two end without a treatment, because the answer is a hand-over. One ends without a fact, because the fact was sought and not established.

19.1 A decompensated heart and a request for an antidepressant

The referral. From a cardiology ward. A man in his sixties, readmitted within months for decompensated heart failure. Low mood, poor appetite, sleeping badly, no motivation, will not participate with physiotherapy. Please see and start an antidepressant.

The reasoning. The failing circulation can generate every item on that list. Fatigue, appetite change, broken sleep, slowed movement and withdrawal from effort follow equally from the haemodynamic state, from the treatment burden, or from a mood disorder, so the count will be high whichever is true and says nothing about which of the three is operating.

What separates them is **order in time**. Put the cardiac events, the regimen changes and the mood changes on one timeline. If low mood arrived with the decompensation and lifted as it settled, that is a person responding to a frightening event. If the breathlessness and the swelling improved while interest, outlook and self-reproach stayed where they were, the mood state has separated from the haemodynamic state and behaves as an illness in its own right. The second lever is pleasure in things costing no physical effort, because a man too breathless to reach the gate can still say whether a visit gives him anything.

Keep the treatment limb open: the regimen is itself a psychological load, and admission interrupts whatever he took at home. Then look again at what was asked. The note has gone straight to a prescription, and agent choice in established cardiac disease is not this chapter's to give: it belongs to the psychopharmacology notes, and prognosis with safe antidepressant choice after myocardial infarction to the merged organ-interface routing section.

TRAP **Accepting the triage the referral has already performed.** The note has decided that this is a psychiatric disorder and that the intervention is a tablet. A consultation answering only the question asked leaves the attribution unmade, and reads later as agreement.

What the examiner is testing. Whether you attribute symptoms rather than count them, and whether you can be useful to a cardiology team without writing a prescription.

19.2 Breathless at night, and already told it is anxiety

The referral. A woman in her forties with obstructive airways disease is referred to the psychiatric outpatient clinic with episodes of breathlessness, palpitations and terror that wake her from sleep. The letter states that her airflow is stable in clinic, that the attacks are anxiety, and asks for treatment of panic.

The reasoning. The whole difficulty sits in the letter's middle clause. Breathlessness and panic recruit the same body and the loop closes both ways: airflow limitation produces fear, fear produces hyperventilation, hyperventilation produces breathlessness. The features shifting the probability, the temporal shape of the episode, its relation to exertion or rest, and the content of the thought at the peak, are taught by this chapter's section on obstructive airways disease, asthma and chronic respiratory failure. This vignette is about what a clinician does when they shift the probability without settling it, which is most of the time.

Two things follow. The commonest correct answer here is both rather than either, and each condition worsens the other, so treating one and closing the case leaves the disability where it was. And the errors are not symmetrical: calling a panic attack an exacerbation costs an unnecessary course of treatment, while calling an exacerbation a panic attack can cost a life. Note too what the letter's premise establishes. Stable airflow in a daytime clinic is a statement about the clinic, not about three in the morning, and nocturnal episodes have nocturnal explanations a daytime measurement cannot exclude.

One contribution is concrete, not diagnostic. Coexistent anxiety and depressive symptoms are routinely used to filter patients out of the rehabilitation that treats the deconditioning, and the guideline position blocking that is stated, in its exact wording, by the owning section.

RULE **When the triage cannot be completed in the room, complete the safe half and record why.** Treat the airway first, interrogate the psychiatry once the episode has settled, and write down which half you answered today. A note committing to a label it cannot defend is more dangerous than one stating what it could not decide.

What the examiner is testing. Whether you accept a triage handed to you in a letter, which direction of error is dangerous, and whether you can offer both with an order of work.

19.3 Confusion after transplantation, and what the psychiatrist does not own

The referral. A call from a transplant unit, late evening. A man in his thirties, some weeks after receiving a kidney from a living relative, has become confused, does not recognise the ward and has had a witnessed seizure. He is on maintenance immunosuppression and has had pulsed treatment for suspected rejection in recent days.

The reasoning. Several sufficient causes are present at once, so the examinable skill is **sequencing rather than selection**. He may be septic while barely febrile, metabolically deranged from graft dysfunction, hyponatraemic or hypoxic, on more than one agent with recognised

neuropsychiatric toxicity, and carrying the load of a graft he believes he could lose. An opinion that names the drug and stops has chosen the interesting cause rather than the first one.

The sequence that survives scrutiny puts reversible and dangerous causes first, keeps the immunosuppressants permanently on the list rather than at the end of it, and treats the psychological formulation as concurrent rather than as a diagnosis of exclusion. One technique is worth taking away whole: the drug history here is a sequence, not a list. What was given, when relative to the event, and what was given alongside it carry the information, because the labelled neurotoxicity signals in this class attach to changes of regimen and to combinations. A chart transcribed as current medicines has discarded the structure that made them reportable.

The boundary, which is the point of this case. Three of the four things the team wants are not yours here. Detection, rating and management of the acute confusional state, including any drug used for it, are taught once, with the stroke, traumatic brain injury, delirium and focal brain syndromes notes. Which co-prescribed agent alters exposure to which immunosuppressant belongs to the psychopharmacology notes. The named neurotoxic syndromes of the individual agents belong to this chapter's transplant and immunosuppressant section. You hand over a dated mental state, a timeline against the drug chart, a statement that the immunosuppressant stays on the differential, and the pathway for each remaining question.

RULE **A referral you cannot answer completely is answered by naming who answers each part.** Routing is a clinical act, not an evasion. It moves a question to somebody who can settle it, and it beats a partial version from memory that commits the receiving team to something you never verified.

What the examiner is testing. Whether you sequence rather than select, and whether you can be useful while declining most of what was asked.

19.4 Two questions from a surgical ward, and three different owners

The referral. The surgical team telephones the day after major abdominal surgery on a man in his seventies. He is agitated at night, pulling at his lines, settled by day. They ask two things in one breath: does he have capacity to refuse the drain, and is this postoperative cognitive dysfunction.

The reasoning. Separate the questions before answering either, because one undifferentiated account of postoperative confusion loses marks in both directions at once. Capacity first. One property does most of its work here: capacity is decision-specific and time-specific. The question is never whether this man has capacity as a standing attribute, but whether he can make this decision at this hour, and the perioperative window supplies an unusual density of moments in which the answer moves: pain, disturbed sleep, premedication and the hours after anaesthesia all sit inside it. An opinion formed on the morning round does not transfer to a conversation that evening. Everything else is elsewhere: the statutory test and the procedure following a finding of incapacity are taught with the mental health legislation notes and are neither restated nor cited by provision. You describe functional ability for the named decision at the named time; the decision belongs to the treating team.

The second question turns on an observation nobody on the ward tonight can make. A cognitive change is a relation between two observations, so the first thing to establish is not what is wrong with him tonight but whether anybody measured him before the operation. If not, the trajectory question cannot be settled on this ward at all, and the honest note says so, manages the acute presentation through the chapter owning it, and arranges the review after discharge.

TRAP **Recording a global capacity verdict.** A note stating that the patient lacks capacity, with no decision named and no time attached, is unusable by the team receiving it and indefensible if later examined. Worse, it travels: written once, it is quoted for months across decisions it was never about.

What the examiner is testing. Whether you separate two questions that arrived as one, and whether your first move is to ask about the measurement taken beforehand.

19.5 A cognitive complaint after treatment ended, and a question the evidence does not answer

The referral. A woman in her fifties, referred by her family doctor some months after finishing chemotherapy. Back at work on reduced hours, she cannot hold several things in mind, loses words she knows well, and works far more slowly than before. She asks whether this will get better, and when.

The reasoning. Take the complaint seriously and decline its title. The word patients use for this names a cause and, in naming it, stops the clinician looking. Run the triage instead.

- The illness limb: fatigue, disturbed sleep, pain and the endocrine consequences of treatment all fail exactly these effortful functions.
- The treatment limb: the cytotoxic agents are one contributor inside a bundle that also holds the surgery and the anaesthesia.
- The independent-disorder limb: an untreated depressive episode degrades attention, retrieval and speed, and is the most treatable thing on the list.
- A fourth possibility peculiar to this population: something still evolving underneath, so whether the picture is static or moving is asked deliberately.

The declared gap, which is her actual question. She has asked for a trajectory. What can be said is that a change of this kind has been reported as still detectable at the end of the interval examined, and that an interval describes how far the looking reached, not where the problem ended. A statement of the course beyond that point would rest on longer follow-up in the oncology and neuropsychology literature, and none was obtained for these notes in a quotable form. The absence describes what these notes opened, is not evidence that the course is unknown to the field, and must not be converted into reassurance or into catastrophe.

What you do without it is concrete. Treat the mood, sleep and pain degrading her cognition today. Characterise the deficit well enough for her workplace to be adjusted around it, whether or not it reverses. And set a review point, because the moment the oncology course closes is the moment nobody is reviewing cognition.

RULE **A declared gap is a complete answer and it has three parts.** State what was sought, on whose authority the missing claim would rest, and what you do without it. An answer carrying only the first two is a shrug; one supplying a remembered figure is worse, because the person receiving it cannot correct it.

What the examiner is testing. Whether you convert an observation window into a prognosis, and whether you can still leave the room with a plan after saying that something is not established.

19.6 Psychosis in an established autoimmune disease, where the differential is the whole answer

The referral. A woman in her twenties with a long-standing systemic autoimmune disease is admitted under medicine with a few days of persecutory beliefs, hallucinations and badly disturbed sleep. Her disease has recently been treated more intensively. The team has sent a serology panel and asks whether this is the disease affecting her brain.

The reasoning. The referral asks for an attribution and offers a laboratory result as the route to it, and neither is available in the form requested. Serology establishes that she has the disease. It does not establish that the syndrome in front of you was produced by it, and a patient with unambiguous systemic disease can equally have an infection, a drug effect or an unrelated psychiatric illness. Attribution here is a reasoned judgement rather than a measurement, and it is the substance of the case rather than its preliminary.

So the answer is an ordered list of what must be displaced before the disease is credited. Infection first, and staying first, in a person immunosuppressed by both the illness and its treatment. Metabolic derangement and blood pressure next, particularly where there is renal involvement. Then the drug chart, read with the dates set against the onset, because the treatment of the disease and the cause of the psychiatric picture can be the same agent; that construct is owned by the routing section of this chapter and by the endocrine and metabolic notes, and is named here only as a direction of travel. Then a primary psychiatric illness with a history predating the systemic diagnosis, and the ordinary response of a young person to a relapsing, visible illness with an unpleasant treatment.

The order matters because of where the answers lead: one path escalates immunosuppression and the other reduces it, so an attribution made the wrong way round is a harm rather than an inelegance. The manifestations and the treatment recommendation with its restrictions are taught in full by this chapter's section on the condition. Give the attribution as a judgement, with the observations it rests on and what would change it. A judgement can be revised without anybody losing face; a finding is very hard to withdraw.

TRAP **Treating a positive result as the answer to the referral question.** The panel speaks to the disease, never to the syndrome, and answering with it skips the contested part of the case. It is also how a commoner and more reversible cause is missed while treatment is escalated in the wrong direction.

What the examiner is testing. Whether you recognise that the difficulty is attribution and not treatment, and whether you write a judgement rather than a verdict.

19.7 A burns unit referral, where the first question is how the injury happened

The referral. A burns unit asks for review of a man in his twenties, several weeks into an admission after a flame injury. He is described as low, withdrawn, sleeping badly and increasingly reluctant to cooperate with dressing changes. The unit would like something started.

The reasoning. The triage runs as it does everywhere in the chapter, but the order is unusual, because one question sits underneath all three limbs and gets harder to ask with every day it is left. Ask first how the injury happened. A burn may be an accident, an assault or a deliberate self-inflicted act, and the risk formulation, the conversation with the family, the safeguarding position and the discharge plan differ completely across those possibilities. Ask early, ask directly, ask the patient and whoever brought him, and record the accounts separately, because they harden with retelling and a unit that has settled on an accident will not re-open it.

Then work the ordinary causes, which this setting stacks. Pain, and specifically the anticipatory dread of a procedure repeated on a schedule, which is a distinct and treatable phenomenon rather than a general lowness. Sleep, degraded by the injury, by the analgesia and by a unit lit and busy through the night. Acute confusional states, which occur here and route out entirely to the notes owning delirium, including their bedside assessment and every drug used for them. And an independent depressive episode, which this setting makes easy to explain away, since comprehensibility is not a diagnostic exclusion.

The word uncooperative is a **referral word**, and it repays unpicking, because it names a behaviour and hides its reason. Reluctance to face dressings after a course of them is a learned response to a procedure that has reliably hurt, and it answers to analgesia, prediction, a share of control and specific procedural anxiety work, none of which is persuasion and none a tablet. Beyond the admission sit appearance, the reactions of others and return to work, all of which belong to this chapter's burns section.

TRAP **Reading a refusal of dressings as either a character problem or a depressive symptom.** Both readings skip the mechanism and both produce the wrong intervention, one punitive and one pharmacological. The behaviour is usually specific to the procedure, which is exactly what makes it treatable.

What the examiner is testing. Whether you ask how the injury happened before anything else, and whether you can decompose an unhelpful referral word into treatable parts.

The cases share a shape: the referral names a symptom and requests a treatment, while the real question is about order in time, the rest of the chart, or ownership.

THE REFERRAL AS WRITTEN	THE QUESTION UNDERNEATH IT	WHO OWNS THE ANSWER
Low mood in heart failure, start an antidepressant	Did mood move with the circulation, or separately	Attribution here; agent choice, psychopharmacology
Her attacks are anxiety, treat panic	What drives tonight's episode, and which error is dangerous	The obstructive airways and respiratory failure section
Confused and fitting after transplantation	What was given, when, and alongside what	Sequence here; the state itself, delirium; interactions, psychopharmacology
Capacity to refuse, and is this cognitive dysfunction	Which decision, at what hour, and was he measured beforehand	Decision-specific and time-specific here; the test, legislation
Will my concentration come back	What degrades it today, and over how long a window was it observed	Function here; the findings, the chemotherapy and interferon section
Is this the disease affecting her brain	What must be displaced before the disease is credited	Attribution here; manifestations, the owning section
He is not cooperating, start something	How did the injury happen, and what is he avoiding	Circumstances here; the sourced material, the burns section

Four moves sit under these cases, and not one of them is a fact to recall.

Re-opened, not accepted	**Attributed, not counted**	**Named, not attempted**	**Sought, not remembered**
THE TRIAGE THE REFERRAL ARRIVES WITH	WHAT MAKES A SYMPTOM LIST MEAN ANYTHING	THE PART OF THE REFERRAL THAT IS NOT YOURS	A FACT THESE NOTES COULD NOT ESTABLISH

The first move is the one the other three depend on, because a triage accepted at the door decides the answer before any of them is run.

IN INDIA

Several features of Indian practice change how these cases run, and none is a figure, because this section holds none. Referrals arrive late and pre-decided, so the first act is usually to re-open a triage somebody else closed. A relative is almost always present, which makes collateral history easy to obtain and the funding question answerable at the bedside. Follow-up happens by travel rather than by schedule, so a plan whose active ingredient is a review appointment fails quietly. And how a burn happened carries legal consequences from the day of admission.

GOOD TO REMEMBER

Before answering any referral in this chapter, ask three questions in order. What has changed, in what order, against what baseline. What else on the chart or in the admission could produce the same picture. And which part of this is not mine to answer here. The first settles the attribution, the second protects the patient, the third makes the note usable.

Run the triage rather than accepting the one in the referral letter, attribute by order in time rather than by symptom count, name the owner of every question that is not yours, and where a fact was not established, say what was sought, whose claim it would be, and what you do anyway.

20 · Consolidation and quick review

In the last week before the examination the useful question is not what this chapter contained but what it installed, and it installed one triage, one habit of naming owners, and one way of saying what is not established. Every section is a worked instance of those three. This consolidation restates them in the smallest form that survives questioning: the triage and the evidence deciding each limb, a grid of the features separating the conditions the chapter teaches, the boundaries stated as a competence rather than an apology, the shapes an answer should take, and what to read if there is an hour left. Read it as a set of moves rather than a summary, because content cannot be compressed without becoming the thin second copy the chapter argued against, whereas reasoning can.

20.1 The three-way question, and what each answer is made of

The chapter opens with a triage and never leaves it. A psychiatric presentation in a medically unwell patient is being produced by the illness acting on the brain, by the treatment for that illness acting on the brain, or by a psychiatric disorder that is the patient's own and has arrived in a body that happens to be unwell. Those are the three limbs, and the chapter is training in choosing between them and in saying why.

Two properties matter more than the categories. The limbs are not exclusive, and more than one is usually running at once, so the working question is never which limb is true but which is doing the most damage now and which, if missed, is dangerous. Each is settled by a different kind of evidence, which is why an answer naming a syndrome without placing it against a limb has skipped the examined step.

LIMB	WHAT SETTLES IT	WHAT IT OBLIGES	CHARACTERISTIC FAILURE
The illness acting on the brain	Disease course and physiology, read against the timing of the psychiatric change	Establish it and hand back for treatment of the disease, since no psychiatric intervention holds while the process continues	Waiting for a confirmatory test in conditions that have none
The treatment acting on the brain	Exposure and interval, from the drug chart and the procedure record rather than the mental state	Raise the exposure with the treating team, and keep the agent permanently on the differential	Pairing an agent with a syndrome and omitting the interval
A psychiatric disorder in a medically ill person	The patient's own history, baseline and account of themselves, with collateral	Treat it in full as a psychiatric disorder, since this limb is the largest and the least delegable	Accepting that anyone would reasonably feel that way, and stopping there

RULE **Name the limb, then name the evidence that placed you in it.** The follow-up at this interface is one question in different words: what else could be producing this, and what would separate them. An answer that has already given the limb and its evidence has answered it in advance, and bought time to be interesting about the rest.

20.2 The manoeuvres that decide it, in the order they are done

The triage is settled by history far more often than by investigation, and by a small set of manoeuvres that cost nothing and are frequently skipped. They are worth more than any single association here, because they transfer to every section and beyond the chapter.

- **Get a baseline from someone who is not the patient.** A medically unwell patient is observed by staff who never met the well version, so the comparison that would make a change visible does not exist on the ward until a psychiatrist goes and gets it.
- **Read the tempo, and read it for mismatch.** A presentation moving faster than the condition it resembles, continuing to move while adequate treatment is given, or following a slope rather than an episode, is telling you something.
- **Lay the psychiatric change against two timelines.** The medical course with its events dated, and the drug chart read for start dates rather than for the current list. This is the single act that separates the first limb from the second.
- **Ask what was stopped as well as what was started.** An admission removes alcohol, nicotine and whatever else was used at home, and often interrupts the patient's own long-term treatment. That syndrome, on a medical ward, is read as a complication of the illness.
- **Take consciousness and attention before anything else in the mental state.** A fluctuating disturbance of attention and awareness changes the entire answer and moves the problem into territory this chapter names and does not teach. Recognising the shape and handing it on is the examinable act.

TRAP **Settling for understandable.** The commonest error at this interface is not a wrong diagnosis but an abandoned one: the low mood is judged a reasonable response to a serious illness and the assessment stops. Comprehensibility is not a diagnostic exclusion and would not be accepted for any other comorbidity. It costs more here, because the psychiatric state is often what removes the patient from the treatment that would have changed the medical trajectory.

20.3 Telling the systemic conditions apart, which is how they are examined

The autoimmune and inflammatory band is not examined by asking what each condition is. It is examined by asking how you would tell them apart with a patient in front of you, which is why the six conditions sit together on one comparator plate. Two axes organise them: where the discriminating information sits, and what kind of claim you hold when you speak.

CONDITION	WHERE THE DISCRIMINATION SITS	THE SEPARATING SENTENCE	WHAT YOU MAY NOT SAY
Neuropsychiatric lupus	In the reasoning, not in any test	Diagnosis is **attribution**, built from timing against disease activity, absence of a competing explanation, and the character of the syndrome	That any investigation confirms it, or that one manifestation is the commonest
Neurosarcoidosis	Inside the mental state you are already examining	A picture crossing the confusional, psychotic and catatonic layers together, and it can present before the systemic diagnosis exists	That the presentation is common or typical, since a presenting feature is not a frequent one
Primary vasculitis of the central nervous system	In a specialty literature this text could not quote	A declared gap: no self-contained quotation explicitly linking it to a psychiatric manifestation was located	Any named psychiatric syndrome, and equally that it has none, since a failed search is not a negative finding
Antiphospholipid syndrome	Outside the mental state, in the record and the serology	The reported psychiatric range reads like a roster of general adult psychiatry, so the mental state is where the discrimination is not	That the antibody caused the syndrome, since the source reports co-occurrence and no proportion
Coeliac disease	Not in a neuropsychiatric sign at all	A co-occurrence claim against healthy controls, spanning neurodevelopmental, affective, anxiety and eating-disorder categories	That it is disease-specific or causal, since the comparison against another chronic illness was never made
Inflammatory bowel disease	Not in a sign either, but in the course over time	A covariation claim with two arrows and a qualifier: effects run both ways and may influence disease course and psychological health	That psychiatric treatment alters the natural history of the bowel disease

Read the third and fourth columns before the first. Two of the six carry no neuropsychiatric sign of their own, and that asymmetry is the point of putting all six on one plate, because the row not behaving like its neighbours is the row an examiner will use. One condition sits adjacent to the family without belonging to it: the paraneoplastic syndromes are immune-mediated, but their defining feature is a sequence rather than a sign or a claim type, since the syndrome arrives before the malignancy is detected, which is why suspicion, and not the workup, is the psychiatrist's part.

RULE **Name the kind of claim before you name its content.** Reported in is not caused by, an association is not a mechanism, a co-occurrence is not a covariation, and a presenting feature is not a frequent one. This distinction carries more marks than any single disease fact, because almost every systemic association here is reported at the level of co-occurrence. An answer keeping the claim and the clinical move apart reads as trained; one fusing them reads as memorised.

20.4 The treatment limb, compressed to exposure and interval

The second limb has one shape wherever it appears. The decisive evidence is in the drug chart and the procedure record, the question is what was given and when relative to the change, and the detail of what to do belongs elsewhere. What you carry is an agent family paired with a syndrome and with an **interval**, meaning the time between that exposure and the psychiatric change.

EXPOSURE	THE QUESTION IT FORCES	WHERE THE DETAIL IS TAUGHT
Calcineurin inhibitors after transplantation	What else was given, in what order, and when relative to the event, since one labelled signal is conditional on a combination	Label level in the transplant section; prescribing and interactions in the psychopharmacology notes
Corticosteroids	Which agent, over what course, and what was the earlier psychiatric history	Routed out of the chapter to the endocrine and metabolic notes, so there is exactly one place to look
Cytotoxic treatment	Did the difficulty predate the first cycle, does it fluctuate within a day, is it evolving, and has it an objective correlate	The cancer-drug section, which teaches the construct rather than the patient's word for it
Interferon-alpha	Prevention before a defined course or treatment after the syndrome declares itself, and in which population	The cancer-drug section; interferon-beta belongs to the neurology and psychiatry interface notes and is never pooled with it
Anaesthesia and surgery	Was anything measured before the operation, since a change is a relation between two observations	The perioperative section for the construct; the delirium notes for the acute confusional presentation
What the admission removed	What was stopped, at what hour, and whether the timing fits	Here, in the triage, since no other chapter owns a removal as an exposure

TRAP **Naming the agent and the syndrome and stopping.** That pairing is the least discriminating half of the topic and every candidate has it. The interval tells a clinician when to look and which explanations compete at that moment, and the population tells them whether the finding applies to this patient. A finding stated without its population is not a shorter answer, it is a weaker and different one.

20.5 The boundaries, which are a competence and not an apology

Liaison work is routing, so a chapter that practises routing is teaching the clinical skill rather than excusing an omission. A referral arrives from another team, written by someone unsure what to ask, and the first act is deciding whose question it is: the treating team's, the psychiatrist's, or a third specialty's. A resident who has practised deciding where a construct belongs has been practising the job; one who memorised a thin paragraph about each has not, and the paragraph fails the follow-up. Stating the boundary aloud also makes an answer checkable: whether a chapter mentioned a topic is a test it can pass while leaving its reader unable to answer anything, whereas whether the reader can name the owner has a right answer.

THE QUESTION IN FRONT OF YOU	WHO OWNS IT
Which agent, at what dose, adjusted how for renal, hepatic, cardiac or transplant disease	The psychopharmacology notes. This chapter teaches everything preceding a prescription and adds no second version of it.
Is this an acute confusional state, and how is it assessed and managed	The stroke, traumatic brain injury, delirium and focal brain syndromes notes, including every bedside instrument, rating scale and drug choice.
Does this patient have capacity, and what is the statutory test	The mental health legislation notes. This chapter says only that capacity is decision-specific and time-specific.
The general immune-to-brain model, and the workup of an autoimmune encephalitis	The neurology and psychiatry interface notes, which also own interferon-beta. This chapter owns the disease-specific presentation and the decision to suspect.
Depression and anxiety after infarction, critical-illness sequelae, dialysis, depression and demoralisation in cancer, palliative care psychiatry	The consultation-liaison, geriatric and transcultural psychiatry notes. These are defined by a setting and a referral pathway, not by an organ.
Psychiatric disturbance following a corticosteroid	The endocrine and metabolic psychiatry notes. This row follows a drug rather than a setting, and follows it out of the chapter.
A possible assault, and the pathway that follows it	The forensic and medicolegal material. Raised in the burns section as an antecedent question and handed on at once.

RULE **Route in the answer, out loud, before you are asked to.** A cross-reference is an address and not an abstract, so it will not survive being revised from, and the failure mode is a candidate who knows a construct matters but cannot say what it does. Say what you are handing on, who holds it, and what you keep. That costs seconds and demonstrates the judgement this interface examines.

20.6 The declared gap, and the sentence that carries it

This chapter declares more gaps than most, deliberately, and the declarations are examinable content rather than housekeeping. A candidate who can say what is not established, and on whose authority a statement would rest, answers better than one who fills the space, because a filled space is unverifiable and frequently wrong. A **declared gap** has three parts, each load-bearing.

State what was sought, with its boundaries: a quotable statement of a particular kind, linking a named condition to a psychiatric manifestation, in a source of stated standing, rather than a vague admission of uncertainty. Name whose literature would hold it, since the specialty following these patients owns the phenotype and deferring to it by name is a different act from not knowing. Then say what you do anyway, because a missing figure never changes the clinical move: prevalence tells you how hard to look, never whether this patient has the syndrome.

The declared gaps are worth revising as a set, since a stem often lands on one. The psychiatric phenotype of primary vasculitis of the central nervous system was not obtained, nor was a specialty statement binding psychiatric management to chronic respiratory failure. No prognostic figure is attached to anxiety after infarction, and no usable recommendation was opened on the desire for hastened death. The epidemiological specifics ordinarily quoted about the stress-precipitated cardiomyopathy are absent, and that section teaches what such a figure would have to measure. In lupus the antibody associations and every frequency statement are absent, because without a diagnostic gold standard each study defines its own numerator. In burns the prevalence figures and the work-reintegration comparison were not obtained, while the Indian work on self-immolation was. And no mechanism is asserted anywhere in the gut-brain material, because neither finding it holds identifies a pathway.

TRAP **Turning a missing number into a negative finding.** The candidate who cannot recall a figure says the association is weak or the condition is rare. That is a categorical claim built on a failure of memory, and it is unsafe where being wrong costs a life. Absence of a located quotation is a fact about a text's sources, never about the disease.

20.7 Answer structures, and the order the reasoning goes in

Almost every question at this interface is one of four stems, and each has an opening move that buys the rest of the answer. Learn the openings, because the first two sentences are where an examiner decides whether you are reasoning or reciting.

- **Discuss the psychiatric aspects of a named medical condition.** Open with the triage applied to that condition, naming which limb it mostly trains and why. Do not open with a list of syndromes that could have been written about any organ.
- **When would you suspect an occult or systemic cause.** Open with the readings that raise suspicion: tempo, the company the syndrome keeps, and whether a general history was taken at all. Then say what you would do, and route the investigation to its owner.
- **How would you tell these conditions apart.** Open by naming them separately and refusing the single undifferentiated paragraph about inflammation reaching the brain. Give each its discriminating feature, and mark which parts rest on a source and which on clinical reasoning.
- **A patient with a known systemic diagnosis presents with a psychiatric syndrome.** Open with the exclusions, in the order infection, metabolic disturbance, the treatment itself, and primary psychiatric illness, because those are commoner and more reversible than the manifestation being reached for.

The body then runs in the same order whichever stem opened it. Characterise the syndrome in the patient's own words. Place it against the three limbs and say what put it there. Give the discriminating feature separating it from its nearest neighbour. State each finding with its population and observation window attached. Name every construct you hand on, with its owner. Declare the gap where the chapter has one, in its three parts. Then say what you would do on the ward tomorrow, which is what most answers run out of time for and what a marker is waiting on.

Two sentences are worth rehearsing aloud, because candidates who have never said them tend to bluff. The gap sentence: that figure is held in the specialty's literature, I cannot state it from a source, and it does not change what I do with this patient. The attribution sentence: this is my judgement rather than a finding, and here is what would change it.

Both sentences are built from the same four fields, and so is every finding this page asks you to state. The comparator grid asked what kind of claim each row holds, the exposure table asked when and in whom, and the gap declaration asked what is not established and whose literature would hold it. Four words carry all of it.

MNEMONIC

Claim, Population, Window, Gap

Claim: say what grade of statement you are holding before you say what it contains, since the grades separated in the comparator rule above are not interchangeable and the weakest of them is the one most often reported here. **Population**: say in whom it was observed, because leaving that out does not shorten the answer, it changes what the answer means. **Window**: give the period over which it was seen, and on the treatment limb the gap between the exposure and the change, which is the half of that pairing most candidates omit. **Gap**: where nothing can be stated, declare it in the three parts set out above. Asked in that order the four survive a condition this chapter never taught, which is why they outlast any single row in the tables.

20.8 If there is an hour left

Revise in this order and stop when the hour is gone rather than when the chapter is finished. The order puts first what an unfamiliar stem is most likely to need.

- **The triage and its evidence table.** It is the answer structure for every question here, it works on constructs you have never met, and it is the shortest thing on this page.
- **The manoeuvres.** Baseline, tempo, the two timelines, what was stopped, and consciousness first.
- **The comparator grid for the six systemic conditions.** Work across each row, ask what makes that condition itself rather than the one above it, and check what kind of claim the row is.
- **The owner list.** Prescribing, delirium, capacity, the immune-brain model with the encephalitis workup, the setting-defined liaison constructs, and the corticosteroid.
- **The gap sentence in its three parts,** and the specific gaps this chapter declares, since a stem landing on one is a stem you answer well while others invent.

- **The exposure and interval table,** read for the second column, because agent to syndrome is the half everybody has and the interval is the half that separates.

What can safely be left is every list this chapter did not build an answer around, and every construct owned elsewhere. Reading a routed topic here in the last hour costs the material only this chapter carries: systemic and inflammatory disease presenting psychiatrically before anybody knows it is present.

If only one card leaves the hour with you, make it this one: four questions to ask of any finding before you say it aloud.

Claim	Population	Window	Gap
WHAT KIND, BEFORE WHAT CONTENT	IN WHOM IT WAS OBSERVED	OVER WHAT PERIOD, AND HOW LONG AFTER	NOT ESTABLISHED, AND WHOSE LITERATURE HOLDS IT

IN INDIA

No India-specific figure is held for this consolidation and none has been invented. The Indian point is about information rather than epidemiology. In most Indian general hospitals the psychiatrist reaches the medical bedside on request, usually late, to a referral question written by a team under pressure, and often without the previous notes, the collateral history or the drug chart's start dates. Every manoeuvre above depends on information the referring team holds and the letter rarely contains, so the skill is asking for it explicitly and early rather than working from the mental state alone. Distance and cost then decide whether the plan is followed at all.

GOOD TO REMEMBER

Three questions at the bedside, one afterwards, and one sentence held in reserve. Is the disease doing this? Is the treatment doing this? Is this a psychiatric disorder in a person who is unwell? Then, before writing: whose chapter answers the part I am handing on? And in reserve, for the moment a figure is wanted and cannot be sourced: that number is held in the specialty's literature, I cannot state it, and here is what I would do regardless.

The chapter installs a triage, a habit of naming owners, and a way of declaring what is not established. Everything else in it is worked examples, and a candidate who holds those three can answer on a condition this chapter never mentioned.

Index

Every entry jumps to its page. Numbers are the page of first mention.

#

10-20 International System 309
3-alpha-androstanediol 811
4A's Test (4AT) 941
5-ht2a 734
5-HT2A antagonists 1268
5-HT2A inverse agonist 1107
5-ht7 antagonism 1755

A

abbreviated application 1749
ABCDEF bundle 960
abducens nerve palsy 1045
Absence seizures 119
absolute neutrophil count 1107
abulia 528
acalculia 999
accelerated long-term forgetting 983
accelerated protocol 1584
acceptance of the diagnosis 240
access 30
access pathway 1797
accumulation 435
aceruloplasminaemia 698
acetazolamide as first-line treatment 364
acetylcholine deficiency 944
acetylcholinesterase 468
ACTH escape 636
activating antidepressant 163
active metabolite 799
active placebo 1716
Active recall 279
activity returns 2016
activity score 1525
Acute akathisia 1303
Acute dystonia 1290
acute encephalopathy 929
Acute intermittent porphyria 703
acute monotherapy 1748
acute myopia and secondary angle-closure glaucoma 367
acute necrotising encephalitis 1178
acute on chronic 1474
Acute post-TBI agitation 925
acute stress disorder 2104
acute suicidal ideation or behavior 1655
acute syphilitic meningitis 1183
acute tryptophan depletion 1707
address route 304
adequate antipsychotic trial 1503
adequate dose and duration 923
adjunct 213
adjunctive therapy 1748
adjunctive therapy to antidepressants 1748
adjunctive therapy with lithium or valproate 1748
adjunctive treatment 745
adjustment 176
administrative counts 1967
adolescent-friendly service 1955
aducanumab 456
adverse drug event 1516
adverse drug reaction 1516
adverse-effect signature 242
afterdischarge 59
age standardisation 1875
agent and jurisdiction cell 1789
agent that does not increase prolactin 1329
AGEP 1613
Agnosia 294
agonist spectrum 1266
agranulocytosis 137
Agranulocytosis 137
Agraphaesthesia 999
agraphia 999
agrypnia excitata 1165
AIDS mania 1089
akathisia 302
akathisia rate 1756
akinesia 987
akinetic mutism 339
Akinetic mutism 339
akinetic-apathetic syndrome 990
albumin-adjusted calcium 685
albuminocytologic dissociation 320
alert, lethargic, stuporous, and comatose 333
alien hand syndrome 1015
alien limb phenomenon 1216
allele frequency 1392
allocation 858
allopregnanolone 811
Allosteric modulation of nicotinic receptors 1417

alpha-1 adrenergic 1260
alpha-1 adrenergic antagonism 1262
alpha-1 adrenergic receptors 1260
alpha-2-delta protein 1416
alpha-synuclein 441
alternative pathway 1277
Alternative psychosis 102
Alzheimer's disease 424
amantadine 752
amnesia 20
amnestic syndrome 447
amstar-2 1707
amyloid cascade hypothesis 442
amyloid PET 506
amyloid plaque 439
amyloid precursor protein 437
amyloid-beta 42 437
amyloid-related imaging abnormalities 599
amyotrophic lateral sclerosis 339
analytical validity 1651
Anosodiaphoria 867
Anosognosia 866
anterior cingulate subcortical circuit 852
anterior pre-optic nucleus 1457
Anterograde amnesia 718
anti-amyloid monoclonal antibodies 1651
anti-NMDA receptor encephalitis 130
Anti-NMDA receptor encephalitis 130
anticholinergic burden 924
anticholinergic delirium 1092
anticholinergic syndrome 1466
anticholinergic toxidrome 1472
anticipation 155
anticipatory monitoring 1632
antimigraine effect of amitriptyline 357
antimuscarinic blockade 1261
antiphospholipid syndrome 2026
antipsychotic polypharmacy 1457
antipsychotic-induced cardiomyopathy 1540
antithyroid autoantibodies 677
Anton syndrome 999
anxiety 15
anxiety sensitivity 2041
Apathy 17
Apathy after traumatic brain injury 913
aphasia 27
APOE 432
apparent clearance 1780
appearance concern 2105
Apraxia 296
arginine vasopressin 636
Argyll-Robertson pupil 1184
ARIA-E 600
ARIA-H 600
aromatic anticonvulsants 246
Arsine gas 823
Arylsulfatase A 706
aspirin-butalbital-caffeine compounds 349
assertive follow-up 1906
associative stigma 1878
astasia-abasia 294
Astereognosis 999
asterixis 302
asymptomatic neurocognitive impairment 1074
AT(N) framework 455
at-risk HIV-negative comparator 1088
atherogenic dyslipidaemia 764
ATP7B 688
attendance management 1962
attention 46
attribution 17
attrition 1135
atypical 50
atypical antipsychotic 469
Auditory agnosia 999
augmentation 662
autoantibody screen 1174
autobiographical memory 1575
autoimmune or paraneoplastic encephalopathy 1170
autoinduction 255
autonomic arousal 29
autosomal dominant 436
avoidance 156
axial features 1130
azapirone 1414

B

Babinski sign 299
Balint syndrome 999
Basal forebrain amnesia 981
Bedside language examination 295
beer drinker's hyponatraemia 780
behavioural and psychological symptoms of dementia 583
behavioural teratogenicity 269
Behavioural treatment 353
benign ethnic neutropenia 1107
benzathine penicillin 1153
benzodiazepine receptor agonists 1407
benzodiazepine-sensitive 1401
bidirectional 15
bidirectional referral 1916
bidirectional relationship 15

bidirectionality 257
bilateral basal forebrain lesions 976
binding profile 1253
Binding promiscuity 1429
Binswanger disease 881
Bioavailability 1272
biogeographic group 1525
biomarker 324
biperiden 389
black box warning 169
Blackfoot disease 823
Bocquet criteria 1619
Bodily activation 376
border crossing 1789
boundary handover 1638
bounded 329
boxed warning 252
Braak staging 441
Bradyphrenia 423
bradyphrenia 423
brain herniation 318
brexpiprazole 589
bridging veins 1029
bromocriptine 756
Bruns syndrome 1190
bulbar palsy 424
bulk drug 1697
burden of normality 205
butterfly glioma 1037
butyrophenones 757
butyrylcholinesterase 576

C

C9ORF72 496
CADASIL 474
CAG repeats 1134
CAG trinucleotide repeat expansion 1135
calcimimetic agents 687
calibration 141
CAM-ICU 940
cancer-related cognitive impairment 2070
capacity-based return 1962
carbamazepine 134
Carbamazepine 134
carboxyhaemoglobin 824
care manager 1941
caregiver support 413
carer organisations 1917
Casal's necklace 727
case definition 433
caseload-based 1941
Caspr2 1165
catalogue answering 2019
Catastrophic reaction 866
catatonia 127
catecholamine surge 2049
central pontine myelinolysis 773
centralised procedure 1655
Cerebral achromatopsia 999
cerebral malaria 1073
Cerebral toxoplasmosis 1083
cerebrospinal fluid 130
cerebrospinal fluid hypocretin measurement 382
cerebrovascular events 429
ceruloplasmin 688
ceruloplasmin oxidase activity 698
CET score 1624
CGRP receptor 360
change control 1742
channelopathy 1236
chemical imbalance 1707
chemobrain 2069
child-pugh class c 1779
chlorpromazine-equivalent 133
cholestatic pruritus 804
cholinergic deficiency 944
cholinergic toxidrome 825
cholinesterase inhibitor 456
cholinesterase inhibitors 456
chorea 302
Chronic akathisia 1305
Chronic amnesic syndrome 979
Chronic daily headache 349
chronic lithium toxicity 1376
chronic low white cell count 1369
Chronic post-TBI aggression 926
Chronic subdural haematoma 1029
Chronic traumatic encephalopathy 566
Chvostek's sign 785
ciclosporin 828
circuit of Papez 967
circuit therapy 1591
circumventricular organs 1209
Clarity-AD 600
Class I antiarrhythmics 1357
class-defining hazard 1436
Clearance 255
clearly defined genetic etiology 1140
clinic-attender sample 2000
Clinical Institute Withdrawal Assessment for Alcohol, revised (CIWA-Ar) 964
clinical utility 1511
clinical validity 1651
clinical-trial no-objection certificates 1685

clonus 84
Clozapine 131
clozapine 131
Clozapine seizure incidence 141
clozapine-refractory schizophrenia 1575
clozapine-resistant schizophrenia 1573
co-occurrence claim 2098
coarse tremor 1369
Cockcroft and Gault equation 807
cognitive and affective items 2034
Cognitive rehabilitation after traumatic brain injury 919
cold retrieval 1248
collaborative care 353
comorbid epilepsy 236
comorbidity 24
company 17
competency 591
competitive employment 1896
complete lactation failure 1559
completed suicide 150
complex partial seizures 31
compound heterozygotes 690
concentration to dose ratio 1530
concussion 845
Confabulation in chronic amnesia 973
confounding 48
Confusion Assessment Method 936
confusional state 77
conjugation to inactive glucuronide metabolites 1394
Consciousness 26
Consolidation 18
constitutive activity 1267
continued use 989
continuity of care 1828
contraction band necrosis 2050
contrecoup injury 893
core delirium workup 953
corneal reflex 294
corrective model 1711
Cortical deafness 999
cortical dementia 296
cortical hyperactivity or excess wake-stabilising orexin 380
cortical-basal ganglia-thalamo-cortical (CBGTC) circuits 1152
Corticobasal degeneration 495
corticobulbar disinhibition 1025
cost of disclosure 1961
cost of illness 1972
cost-effectiveness 960
counter-regulatory hormones 761
coup injury 893
covariation claim 2098
coverage 243
covert non-adherence 1504
creatine kinase 953
creatinine clearance 807
creatinine creep 1545
Creutzfeldt-Jakob disease 312
crude rates 1875
Cryptococcal meningitis 1075
ctri registration 1660
cumulative activity figure 1936
Cushing disease 642
cyclodextrin sodium 809
cyclopyrrolone 1402
CYP2D6 1023
CYP2D6*10 allele 1359
CYP3A subfamily rather than for CYP3A4 alone 1278
cyproheptadine 1453
Cystathionine beta-synthase 708
cysticercotic encephalitis 1192
Cytomegalovirus encephalitis 1096

D

d-penicillamine 701
d3-preferring 1748
Daily Rating of Severity of Problems 812
dantrolene 1458
DAT scan 560
data leakage 1726
dataset shift 1725
date of approval 1663
date of issue 1684
Dawson fingers 1203
day care 1818
daytime extension of hyperarousal 375
Daytime hyperarousal 377
de novo postoperative depression 202
de novo psychiatric disorders 201
de novo psychosis 203
de-esterification 1423
de-induction 1281
Death by neurological criteria 338
decay 1327
Decerebrate posturing 334
dechallenge 1518
decision-specific and time-specific 2027
declared gap 337
Decorticate posturing 334
decriminalisation 1814
deep brain stimulation 1114

Deep brain stimulation 1114
deep TMS 1575
dehydroepiandrosterone 634
deinstitutionalisation 1894
deiodinase 673
Delayed onset of action 1340
delegated gap 1703
Delirium 52
delirium 52
Delirium tremens 944
delivery-side distance 2014
Delusional disorder 852
Dementia pugilistica 921
dementia with Lewy bodies 424
dementia workup 312
Demethylation 1355
denominator 96
Depressive Disorder due to Another Medical Condition 348
Depressive Disorder Due to Another Medical Condition 348
depressive pseudodementia 510
Descemet membrane 695
Dexmedetomidine 951
Dextromethorphan with quinidine 1027
dextromethorphan-bupropion 1495
diagnostic boundary 611
Diagnostic humility 407
Diagnostic overshadowing 405
diagnostic overshadowing 405
dialysis disequilibrium 798
diaphoresis 825
Diencephalic amnesia 965
difficult-to-treat depression 1486
diffuse axonal shearing 899
Diffuse slowing 52
digital divide 1984
digital placebo 1745
DiHS 1620
diphenylbutylpiperidines 1289
disability 165
disconnection syndrome 706
Discontinuation symptoms 813
discrimination 18
discriminator 79
disease-modifying 456
disfigurement 2105
disinhibited syndrome 989
disinhibition 326
displacement 318
disproportionality analysis 1522
dissociation 22
Dissociation 22
dissociative seizure 208
dissociative seizures 208
dissociative symptoms 1500
district mental health programme 1806
Dominant temporal-lobe dysfunction 1001
donanemab 456
donepezil 469
dopamine agonist 756
dopamine agonists 756
dopamine antagonist 468
dopamine dysregulation syndrome 1106
dopamine serotonin antagonist 1293
dopamine transporter inhibition 1343
dopaminergic excess 944
dorsal striatum 1106
dose 17
dose-dependent dual action 1349
Down syndrome 431
DRESS 17
Drug-precipitated REM sleep behaviour disorder 389
Drug-resistant epilepsy 46
DSM-5 65
dual noradrenaline and dopamine reuptake inhibitor 1337
dual orexin receptor antagonists 1651
duration of illness 239
Dysdiadochokinesia 325
dysexecutive syndrome 918
dysgeusia 1402
dyskinesias 130

E

earliest approval-status date in the record 1657
economic evaluation 1530
edaravone 1224
Efavirenz 1099
effect change 1268
effect duration 1422
Efficacy 115
efficacy belief 1889
ELECT-TDCS 1575
electrocerebral silence 340
Electroconvulsive therapy 120
electroconvulsive therapy 120
Electroconvulsive therapy after TBI 925
elemental lithium 1371
elimination half-life 264
emergency act 1444
emergency CSF indication 318
Emotional incontinence 851

emotional lability 202
end-of-life care 589
endpoint selection 1886
engagement 115
environments people live through 1989
enzyme inducer 164
Enzyme inhibition 255
enzyme-inducing AEDs 255
enzyme-inducing anticonvulsant 1395
epilepsy patients 96
Epilepsy-to-psychosis latency 117
epileptic dementia 189
epileptiform discharge 50
epileptogenesis 62
episodic dyscontrol syndrome 184
era mismatch 1679
erythrocyte transketolase activity 713
escalation route 573
ester prodrug formulated as an aqueous nanosuspension 1319
ethosuximide 118
euthyroid 660
evening and early sleep period 377
expectancy 631
explanatory-model assessment 1999
exposure change 391
external validation 1725
Extinction 300
extracellular, synaptic 1165
extrapyramidal nervous system 1295
extrapyramidal side effects 392
extreme delta brush 1160

F

fabricated citation 1484
face-hand test 1228
false neurotransmitters 789
false sense of presence 1104
Familial vascular dementia 886
familiarity heuristic 1436
fasciculations 294
Fast-frequency EEG components 376
fast-off property 1297
fatigable weakness 1232
FDG-PET 461
feeding and reward 381
fetal valproate syndrome 269
fetor hepaticus 789
fidelity 28
Filipin staining 707
finger agnosia 999
first generation 1337
first-generation antipsychotics 131
First-order kinetics 1272
first-pass metabolism 803
first-response knowledge 1879
fit 15
fixed sequence 214
flow 244
fluctuating cognition 457
Fluctuation 287
fludrocortisone acetate 655
fluency illusion 1248
flumazenil 789
flunarizine for migraine prophylaxis 358
Focal aware seizure 20
Focal impaired-awareness seizure 25
focal impaired-awareness seizures 25
Focal onset seizures 21
folate insufficiency 720
forced normalisation 69
Forced normalisation (Landolt phenomenon) 117
Forced thinking 77
Four Questions of Neurology 306
four-field status claim 1647
framework 15
free fraction 806
free recall 464
free-text search 1700
Friedreich ataxia 1218
frontal and basal ganglia lesions 852
frontal beta activity 314
frontal lobe epilepsy 16
Frontal release reflexes 328
Frontal release signs 339
frontal-executive dysfunction 473
fronto-ponto-cerebellar pathway 1023
frontotemporal degeneration 485
frontotemporal dementia 424
full agonist 1266
full gradual escalation regimen 1399
full penetrance 1136
fulminant hepatic failure 702
functional coverage 1801
functional dystonia 1228
functional neurological disorder 207
functional neurological symptom disorder 209
functional or dissociative seizures 212
functional outcome 241
functional status 429
functional unblinding 1715
futility analysis 602

G

galantamine 469
Galveston Orientation and Amnesia Test 916
Gastaut-Geschwind syndrome 177
gastrointestinal hypomotility 1317
gatekeeper 287
general paresis 515
generalised anxiety disorder 156
generalised bullous fixed drug eruption 1623
generalised erythematous rash 1391
generalised HPA-axis activation 377
Generalised tonic-clonic seizures 94
Generalized onset seizures 21
generative model 1738
Gerstmann syndrome 1007
give-way 294
Glasgow Coma Scale 335
glatiramer acetate 1207
Glioblastoma 1037
global mental health movement 1989
globus pallidus interna 1128
glucocorticoid receptor 637
glucocorticoid resistance 1209
glucocorticoid withdrawal syndrome 648
glucuronidation 255
GluN1 (NR1) subunit 1158
Glycogen synthase kinase 3 1370
glycosylated haemoglobin 738
glymphatic system 1209
gp120 1082
graded 137
graded dose-response 1284
gradient 46
grasp reflex 294
greater high-frequency EEG power 376
growth suppression 1561
guided or blended intervention 1984
gut-brain axis 2095

H

halfway home 1896
handover 292
handover boundary 1257
hap 14
hard neurological sign 325
Hardy-Weinberg equilibrium 1526
harm reduction 2002
Hashimoto encephalopathy 659
hatchet face 1236
healer training 1991
hedonistic homeostatic dysregulation 1119
Hemispatial neglect 870
hepatic biotransformation 1626
hepatic encephalopathy 515
hepatic metabolising capacity 802
Hereditary coproporphyria 705
Herpes simplex encephalitis 32
HIFU 1601
high-risk group 168
high-tyramine meal 1361
highest-yield discriminators 215
hip abductor sign 1226
Hippocampal amnesia 976
histaminergic H1 blockade 1261
histone deacetylase 268
HIV-associated dementia 1075
HIV-associated neurocognitive disorder 1073
HIV-associated psychosis 1090
HLA-A*31:01 1527
HLA-B*15:02 1527
hoarding 1119
hobbyism 1119
Hoover sign 301
Hoover's sign 1226
Hospital Elder Life Program (HELP) 959
hub-and-spoke model 1934
hubs 1934
human leukocyte antigen (HLA) variant B*15:02 1391
Humanitarian Device Exemption 1592
Hunter Serotonin Toxicity Criteria 1452
huntingtin (HTT) gene 1134
Huntington disease 423
hydrocephalus ex vacuo 528
Hyperactive delirium 934
hyperaroused and anxious in the daytime 379
hyperbolic 1402
hypergraphia 177
hyperkinetic 37
Hypermetamorphosis 999
Hyperorality 424
hyperphosphorylated tau 444
hyperreflexia 875
hypertonicity 1450
hypnagogic and hypnopompic 384
Hypoactive delirium 934
hypoalbuminaemia 1087
hypokalaemia 366
hypokinetic 1155
hypomagnesaemia 684
hypomania/bipolar-spectrum conditions 344
hyponatraemia 653

I

iatrogenic differential 1637
iatrogenic neuropsychiatric syndrome 1070
iatrogenic phenocopy 1629
ICHD-3 chronic tension-type headache 352
ICHD-3 frequent episodic tension-type headache 351
ICHD-3 migraine without aura criteria 342
ictal 15
ictal affective phenomena 70
ictal depression 70
ictal fear 20
ictal psychic phenomena 20
ictal psychosis 35
ictal symptoms 69
Ictal violence 184
ictus 65
Ideational apraxia 1011
Ideomotor apraxia 999
Idiopathic intracranial hypertension 319
idiosyncratic 246
imidazopyridine 1402
immunotherapy 130
impaired fasting glucose 746
Implicit trauma memory 905
impulse control disorder 817
impulse control disorders 817
inability to switch off arousal-related circuits 378
inadequate response 1479
Incidence 15
index substrate 1533
India ink stain 1094
india's central occupational-safety statute 1948
india's national mental health programme 1975
indicated 118
indigenous resources 1915
indirect sympathomimetic 1363
individual placement and support 1896
indoleamine 2,3-dioxygenase 1211
induction phase 1498
inferior petrosal sinus sampling 649
inflammatory limb 2088
Information Component 1522
inhibition of dopamine uptake 1337
inhibits neuronal sodium channels 1393
innovator application 1749
Inositol monophosphatase 1370
insecticides act, 1968 1866
Insight 288
institutionalism 1907
instrument blindness 1703
insulin omission 761
insulin resistance 733
insurance parity 1830
integrase strand transfer inhibitors 1103
integration 208
intended use 573
intention or postural tremor 1383
intentional non-adherence 2043
inter-ictal psychosis 1550
interface 14
interface of the gamma-2 and alpha subunits 1400
interferon alfa 2072
interferon-alpha 1206
interferon-beta 1206
interictal 15
interictal dysphoric disorder 65
interictal period 65
Interictal personality change 18
interictal violence 184
interleaved retrieval 1435
interleaving 279
intermanual conflict 1015
intermediate allele 1137
intermediate metaboliser 1279
internal inconsistency 558
internal migration 1991
internally displaced people 1965
interpreter 39
intersection route 304
intersectoral action 1993
interval 15
interval between onset and effective treatment 1997
intervention package 1717
intracellular 439
intranuclear eosinophilic inclusions, the Cowdry bodies 1198
Intravenous haloperidol QTc thresholds 957
intrinsic activity 1431
intrinsic factor 720
inverse agonist 1107
ionised calcium 684
ionotropic receptor 1263
irdai 1830
irreversible inhibitors 1363
isolated psychiatric presentation 1175
isolation aphasia 801

J

Japanese encephalitis 1073
JC virus 320
JC virus CSF testing 324

job tenure 1904
Juvenile-onset Huntington disease 1151

K

Kayser-Fleischer ring 690
Kayser-Fleischer rings 695
kindling 16
Kluver-Bucy syndrome 999
Korsakoff syndrome 533

L

L-methylfolate 725
L-thyroxine 669
la belle indifference 302
label-driven prescribing 1637
Lactulose 515
Lacunar dementia 884
lamotrigine 134
lamotrigine half-life 1395
large language model 1725
laryngospasm 1500
lassitude 730
latency 113
lateral hypothalamic neurons 382
law of mass action 1297
layered response 1947
lead lines 819
Lead palsy 822
lead-pipe 294
lecanemab 456
Left lateral prefrontal involvement 904
left-sided limb apraxia 1016
Leipzig criteria 700
lenticular nucleus 690
leucopenia 1392
leukoencephalopathy 324
Lewy body 463
Lewy body dementia 547
LGI1 773
licarbazepine 773
Lifetime prevalence 44
light-near dissociation 1184
Limb-kinetic apraxia 1011
limbic encephalitis 511
limbic system 14
limitation of use 1748
limitations of use 1655
limited awareness, sociocultural beliefs, values and stigma 1842
linezolid 1449
liothyronine 671
Liraglutide 749
loading doses 1509
Localisation 14
localisation-related epilepsy 191
locally meaningful idioms of distress 1991
Locked-in syndrome 337
logopenic variant 487
long QT syndrome 1564
low-grade systemic inflammation 760
lower efficacy than beta blockers 357
lucid interval 65
Lumbar puncture 317
lymphocytic pleocytosis 1161

M

macrocytic anaemia 721
magnetic gait 1037
magnetic seizure therapy 1579
maintenance monotherapy 1748
major adverse cardiac events 1577
major congenital malformations 1554
major depression 142
major depressive episode 105
major neurocognitive disorder 419
Major Neurocognitive Disorder Due to Parkinson's disease 1123
malignant catatonia 663
malignant hyperthermia 1466
mammillary bodies 58
Mammillary-body and thalamic atrophy 971
manganism 820
marketing approval 1655
marketing authorisation 1507
marketing authorisation holder 1507
marketing status 1662
maternal immune activation 1213
Maudsley Staging Method 1484
may be autoimmune 383
may not show increased objective daytime sleepiness 376
McDonald criteria 1202
measured absence 1648
measurement proxy 2017
mechanistic claim 2098
mechanistic routes 1070
media guideline 1864
medical device 1590
Medical Research Council scale 298
medically unexplained presentations 227
Medication-induced parkinsonism 828
Medication-overuse headache 348
medium spiny projection neurons 1136
MELAS 887

melatonin receptor agonist 1352
melatonin receptor agonists 1402
memantine 456
membrane potential 1586
membrane stabilisation 1357
memory monitoring 978
Meningioma 1037
meningovascular syphilis 1183
mental health literacy 1845
mental-status examination 2057
mesocortical dopamine pathway 1296
mesolimbic dopamine pathway 1295
metabolic ratio 1312
metabolic syndrome 632
Metabolic syndrome 632
metabolised mainly via cyp3a4 1755
metaboliser phenotype 1279
Metachromatic leukodystrophy 706
Metformin 720
methyl mercury 823
Methylenetetrahydrofolate reductase 709
metyrapone 650
mhgap 1923
microglial nodules 1083
microspheres 1318
microvascular dysfunction 2050
midbrain panda sign 699
migraine with aura 342
Migraine-depression comorbidity 343
mild cognitive impairment 191
mild neurocognitive disorder 419
Mild traumatic brain injury 337
mild-complicated traumatic brain injury 892
mineralocorticoid receptor 637
Mini-Mental State Examination 554
minimal hepatic encephalopathy 788
Minimally conscious state 332
minimally important difference 1703
Minor physical anomalies 327
mixed dementia 474
mixed features 101
moclobemide 131
mode of action 1253
model of care 1916
Modified Dandy criteria 363
molecular mimicry 821
monitoring proxy 1636
monoamine oxidase inhibitor 389
monohydroxy derivative 1390
monotherapy 108
Montreal Cognitive Assessment 555
Montreal Cognitive Assessment (MoCA) 997
mood disorder due to epilepsy 101
more adverse effects than propranolol 357
Morvan syndrome 1165
Motor impersistence 998
MRI with MR venography 363
MS-induced euphoria 1204
Multi-infarct dementia 474
Multicenter AIDS Cohort Study 1088
Multicomponent non-pharmacological prevention 959
multimodal 567
multinucleated giant cells 1083
Multiple Endocrine Neoplasia Syndrome 656
Multiple system atrophy 1216
multisectoral 1859
muscarinic agonist 1687
muscarinic antagonist 1687
muscarinic cholinergic affinity 1343
muscarinic cholinergic receptors 1260
muscle-specific kinase (MuSK) 1233
must not be used for migraine prophylaxis during pregnancy 359
myocardial stunning 2050
Myocarditis 137
myoclonus 84
Myotonia 1236
myxoedema 659
myxoedematous madness 661

N

Na+/I- symporter 667
narrow therapeutic index 807
national crime records bureau, the ncrb 1871
national health mission 1845
national mental health programme 1803
national mental health survey 1822
national suicide prevention strategy 1852
need-side distance 2014
need-to-service distance 2013
negative neurological finding 403
negative predictive value 670
neonatal adaptation difficulties 1557
neonatal bleeding diathesis 1555
nephrogenic diabetes insipidus 773
nephrotoxic 803
network meta-analysis 1327
neural tube defect 268
neural tube defects 268
neuritic plaque 442
neuroactive steroid 1662
neurocognitive impairment 799
Neurocysticercosis 1073

neurofibrillary tangle 439
neurofilament light chain 505
Neuroglycopenic 762
neuroleptic sensitivity 461
Neuroleptic-induced dysphoria 1329
Neurological Evaluation Scale 330
neurological lens 292
Neurological soft signs 325
neuromuscular excitability 683
neuromuscular hyperreactivity 1466
neuropsychiatry of systemic and neurological disease 320
neuropsychological assessment 223
neurosarcoidosis 320
Neuroscience-based Nomenclature 1293
new-onset mania 1090
NIA-AA 570
NICE migraine-prevention options 355
nicotinamide 728
nicotinic acetylcholine receptor 1232
Niemann-Pick disease type C 707
nigrostriatal dopamine pathway 1295
NINDS-AIREN 477
NLRP3 inflammasome 1209
NMDA receptor antagonism 1260
NMDA receptor antagonist 1027
NMDAR hypofunction 1158
no cdsco approval record was located 1660
nocturnal polyuria 782
non-competitive nmda receptor antagonist 1673
non-confrontational stance 231
non-inferior 1575
non-inferiority 1575
non-invasive brain stimulation 1585
non-response 676
Nonconvulsive status epilepticus 77
nonconvulsive status epilepticus 77
nonspecific symptoms 385
Nontreponemal tests 1188
nootropic 1420
noradrenergic and specific serotonergic antidepressant 1350
norclozapine 1312
Norfluoxetine 1343
normal or only nonspecifically abnormal 309
normal pressure hydrocephalus 525
Normal-pressure hydrocephalus 320
norms of journalistic conduct, 2020 1866
NOTCH3 on chromosome 19 886
npsle 2084
number needed to test 1611

O

obsessive symptomatology 1132
occupancy threshold 1297
occupational health function 1963
Ocular apraxia 999
oculomasticatory myorhythmia 1199
off-target action 1254
offset 736
oil vehicle 1319
old-age psychiatry 1029
oligoclonal bands 320
oligoclonal IgG 320
oligodendroglia 1083
on-off fluctuations 1129
onconeural 1165
onconeural antibody 2078
one-year rule 424
Opening pressure 319
Optic ataxia 999
optimism 239
order in time 2111
Orexin 368
orexin receptor antagonist 1402
orexin receptor antagonists 1402
ORGANIC 23
Organicity 55
Orientation 16
Orientation to person 288
Orientation to place 288
Orientation to situation 288
Orientation to time 88
Orlistat 751
Orobuccal apraxia 1011
orthostatic hypotension 469
osmotic demyelination syndrome 783
osteomalacia 730
ovarian teratoma 130
Overt Aggression Scale 912
overt hypothyroidism 659

P

Paced Auditory Serial Addition Test 1205
paired helical filaments 439
PANDAS 1152
Papilloedema 317
parallel formulation 400
parameter recovery 1738
paraneoplastic 533
paraneoplastic limbic encephalitis 533
parathyroid glands 681
parathyroid hormone 681

parenchymatous syphilis 1183
Parkinson disease dementia 424
Parkinson disease psychosis 1104
Parkinson-plus diseases 1214
paroxysmal EEG finding 311
partial 5-HT1A agonist 1414
partial agonist 589
partial response 729
participation 241
patch testing 1624
Pathological Laughter and Crying Scale 1022
pathway to care 1985
patients with epilepsy 119
Pellagra 711
pellagrinous dementia 727
pentamer 1400
peri-ictal axis 16
peri-ictal category 66
peri-ictal depression 103
peri-ictal period 65
peri-ictal psychosis 68
peri-ictal temporal axis 15
periodic limb movements 373
periodic sharp-wave complexes 312
perioperative neurocognitive disorders 2064
pernicious anaemia 677
Perseveration 430
Persistent delirium symptoms 962
persistent pulmonary hypertension of the newborn 1557
persistent vegetative state 294
perspective 337
phantom shock 2037
pharmacodynamic 133
pharmacodynamic interaction 582
pharmacodynamic interactions 582
Pharmacodynamics 1283
pharmacokinetic 133
pharmacokinetic booster 1023
pharmacokinetic enhancer 1099
pharmacokinetic interaction 140
pharmacokinetic interactions 164
pharmacological discriminator 1428
pharmacovigilance 165
phasic signalling 1419
phenobarbital 70
phenoconversion 1430
phenothiazine 139
Phineas Gage 910
phospho-tau 568
physical dependence 447
physostigmine 1473
pimavanserin 469
piperazine phenothiazines 1289
piperidine phenothiazines 1289
plasma corticotropin immunoassay 654
plasma p-tau217 568
Plasmodium falciparum 1190
platform intervention 1953
pleocytosis 319
Point prevalence 44
police-recorded suicide cases 1871
polycyclic aromatic compounds 1281
polycyclic aromatic hydrocarbons 1281
polyglutamine 1135
poor metaboliser 1277
poor metabolisers 1277
population lever 1864
population register 1941
population-level intervention 2011
porphobilinogen deaminase 703
portable record 1966
positive allosteric modulation 1663
positive allosteric modulator 811
positive control 1111
positive controls 1650
positive inconsistency 301
positive neurological finding 403
positive predictive value 1611
positive signs 302
positive symptoms 113
Post-concussion headache 898
post-concussion syndrome 892
post-ictal psychosis 1540
Post-injection delirium sedation syndrome 1320
post-stroke anxiety 859
post-stroke apathy 859
Post-stroke dementia 874
Post-stroke depression 850
Post-stroke mania 851
Post-stroke psychosis 866
Post-TBI anxiety disorders 904
Post-TBI apathy 928
Post-TBI depression pharmacotherapy 924
Post-TBI depressive disorder 903
Post-TBI disinhibition 926
Post-traumatic amnesia 888
post-traumatic confusional state 917
post-traumatic dysexecutive syndrome 918
Post-traumatic epilepsy 1047
Post-traumatic personality change 909
Post-traumatic stress disorder after TBI 905
posterior dominant rhythm 52
posterior fossa 1037

posterior reversible encephalopathy syndrome 2058
postictal anxiety 100
postictal confusion 16
postictal delirium 88
postictal depression 70
postictal mania 100
postictal psychosis 65
postictal recovery 89
Potency 131
Pramipexole 1115
Pre-lumbar-puncture imaging 318
pre-test and post-test counselling 1144
predictive 294
Predictive testing 1134
Pregnancy Prevention Programme 270
pregnancy prevention programme 270
preictal prodrome 16
preimplantation genetic testing 1145
premanifest testing 1144
premarket approval 1592
Premenstrual exacerbation 812
Premenstrual syndrome 812
Prenatal diagnosis 1145
preparedness 453
prescription digital therapeutic 1725
presenilin 432
presystemic extraction 1276
Priapism 1260
Primary brain injury 893
primary central nervous system vasculitis 2087
Primary CNS lymphoma 1083
Primary hyperparathyroidism 682
Primary polydipsia 780
primary prevention 752
primary progressive aphasia 485
prion disease 515
problem representation 395
procedural act 1444
proconvulsive 1540
prodrug 1319
Progressive multifocal leukoencephalopathy 324
progressive multifocal leukoencephalopathy (PML) 1200
Progressive supranuclear palsy 390
prolactin 212
prolactin neutral 1329
prolactinoma 754
propranolol overdose risk 356
Prosopagnosia 999
protein kinase c 821
protracted situations 1965
protracted withdrawal syndrome 1400
provoked confabulation 973
proxy label 1725
pseudo-Cushing states 650
Pseudo-irreversible inhibitors 1417
Pseudobulbar affect 297
pseudobulbar affect 297
Pseudobulbar palsy 424
pseudodepressed syndrome 989
Pseudohyponatraemia 773
Pseudoparkinsonism 1303
pseudoresistance 1503
pseudosclerotic form 693
psychiatric emergency 24
psycho-pharmacogenomic studies 1530
psychogenic non-epileptic seizure 40
psychogenic non-epileptic seizures 40
Psychogenic nonepileptic seizures 1048
psychological first aid 1945
Psychoneuroimmunology 1201
psychosocial hazards 1960
psychosurgery 1596
Psychotic disorder due to TBI 906
public stigma 1877
public-health implementation 1797
pulmonary rehabilitation 2042
punding 1116
pupillary light reflex 296
pure epileptic amnesia 983
pyrazolopyrimidine 1402

Q

QT interval 683
QTc 786
quantal dose-response 1284
quetiapine 131
quinolinic acid 1211

R

racemic ketamine 1495
Raised intracranial pressure 323
randomised withdrawal 1499
rapid needs assessment 1946
rapid symptom evolution 1172
rapid-acting antidepressant 1651
rapidly progressive dementia 532
reactance 1885
receptor binding profile 1253
Receptor cleanliness 1341
rechallenge 467
recognition 36

recognition capacity 1997
recognition signatures 1237
recovery 16
Red cell folate 723
red flags 431
reduced penetrance 1136
reference schedule 1397
reference standard 212
reference year 1799
referral word 2116
RegiSCAR 1613
regulatory recency 1648
relative infant dose 1558
release 120
release capacity 1707
release mechanism 1423
reliably excludes toxic-metabolic encephalopathy 313
REM atonia 372
REM rebound 389
REM sleep behaviour disorder 373
REM sleep phenomena intruding into wakefulness 384
REM sleep without atonia 388
Remission 44
remote consultation 1820
renal elimination 1531
renal prostaglandins 1378
repeat-length categories 1142
requires trust, rapport-building and open dialogue 1991
research nosology 1074
Reset osmostat hyponatraemia 784
responder care 1951
Restless legs syndrome 387
retained insight 1105
reticular activating system 333
retrieval card 1245
retrieval practice 1244
retrieval template 1245
Retrograde amnesia 916
retrograde amnesia 916
reverse causation 48
reverse inference 74
reverse retrieval 1636
reversible dementia 509
reversible inhibitor of MAO-A 1364
rifaximin 789
right basotemporal-orbitofrontal circuit 861
right lateral prefrontal cortex 852
Rilpivirine 1100
riluzole 1224
ring-enhancing lesions 1083
risk allele 432
risk evaluation and mitigation strategy 1663
risperidone 131
risus sardonicus 693
ritalinic acid 1422
Ritonavir 1099
rivastigmine 469
Romberg's test 301
Rotterdam criteria 815
route knowledge 1889
routed construct 2025
routine CSF profile 319
routine EEG screen is unwarranted 310
Routine pharmacological prophylaxis 961
routing 1653
ryanodine receptor 1458

S

safeguarding floor 2003
safety escalation 1987
safety netting 407
sanctioned coverage 1801
sanguinine 1418
Schedule III controlled substance 1500
schedule iv controlled substance 1781
second generation 1288
Second impact syndrome 891
second-generation antipsychotics 131
secondary adrenal insufficiency 652
secondary amines 1355
Secondary brain injury 894
secondary mania 663
Secondary mania after TBI 907
secondary neuropsychiatric syndrome 1067
secondary polydipsia 780
secondary prevention 479
seizure phobia 154
seizure threshold 47
selective 166
selective loss of hypocretin-producing neurons 383
selective noradrenaline reuptake inhibitor 1425
self-guided intervention 1984
self-stigma 1877
semantic variant 487
semiology 16
Sensory integration 325
separating axis 1434
sequencing rather than selection 2061
Sergei Korsakoff 965
serious 26

serotonergic hypothesis 1023
serotonergic limb 346
serotonin syndrome 346
serotonin toxicity 1341
serotonin transporter 774
serum 25-hydroxyvitamin D 730
serum ceruloplasmin 696
serum homocysteine 723
serum methylmalonic acid 723
service architecture 1797
service variable 1907
set point 44
severe hypertriglyceridaemia 738
severe neutropenia 1107
severe traumatic brain injury 337
sexual dysfunction 182
sham 226
sham comparator 1606
shared method 844
shared-device problem 1985
sharp wave 51
shock-related anxiety 2036
sickness behaviours 1210
sign-symptom mismatch 396
signal 27
SILENT 29
silent antagonist 1266
simple branched-chain fatty acid 1381
simulated treatment 1573
Simultanagnosia 999
simultaneous positive and negative visual phenomena 344
single course of treatment 1663
single-centre trial 1575
Sleep paralysis 384
Sleep-related hallucinations 384
sleep-wake phenotype 399
slow-dose stimulants 1424
slow-release 1351
small-vessel vascular cognitive impairment 883
social determinant 877
sodium and hydrogen exchanger 3, or NHE3 1371
sodium, potassium and chloride cotransporter NKCC2 1377
software as a medical device 1742
source memory 976
spaced retrieval 1248
special population 1395
spectrum bias 1725
spike 33
spina bifida 268
spinocerebellar ataxia type 2 1219
spinocerebellar ataxia type 3 1219
spinocerebellar ataxias 1218
spokes 1934
spontaneous or fantastic confabulation 973
Sports concussion grading 900
spreading neuronal depression 345
stabilisers 242
staccato speech 1472
Stages of Reproductive Aging Workshop 813
staging models 1479
standardised mortality ratios 770
Stanford Neuromodulation Therapy 1575
STAR*D trial 672
state examination 111
state-dependent assessment 403
status epilepticus 20
status true as of 1647
statutory entitlement 1813
steady state 257
stepped care 1921
stepwise or fluctuating decline 876
stereognosis 327
Sternbach criteria 1452
Stevens-Johnson syndrome 246
stock 132
strategic-infarct vascular cognitive syndrome 883
stroke-like episodes 678
strong cyp3a inhibitor 1665
strongly antimuscarinic 1261
structural adaptation 1930
structural measures 1883
structural stigma 1878
structured product label 1655
sub-convulsive 1575
subacute encephalopathy 1094
Subacute sclerosing panencephalitis (SSPE) 1196
subanaesthetic 1496
subclinical hypothyroidism 663
subcoeruleus area 387
subcortical dementia 423
subdural space 1029
subgenual anterior cingulate cortex 1605
subject expert committee 1655
subject expert committee minutes 1655
subjective cognitive decline 538
sublingual tablet 1748
substantia nigra pars compacta 1105
substituted benzamides 658
subthalamic nucleus 1128
subthreshold intensity 1607
sudden cardiac death 1540
suicide surveillance 1852

summary of product characteristics 741
supersensitivity 781
supervised learning 1725
supervision 133
supported employment 1818
supportive source 1648
suprachiasmatic nucleus 371
surface adaptation 1930
surrogate marker 1703
sustained response 1217
Sydenham chorea 1149
symptom-rhythm concordance 2036
symptomatic 45
symptomatic limb 2088
synaptic and extrasynaptic gaba-a receptors 1664
syndrome formulation 402
syndrome frame 398
syndrome of irreversible lithium-effectuated neurotoxicity 1379
Synkinesia, or motor overflow 329
synucleinopathy 389
systematic pharmacology 1252

T

tabes dorsalis 1183
tACS 1600
Taenia solium 1190
Tardive akathisia 1305
tardive dyskinesia 302
target engagement 597
targeted fluid restriction 782
targeted investigation 418
task-sharing 1923
task-shifting 1805
Tat 15
TBI-induced depression 912
TBI-sensitive prescribing 922
tea party epilepsy 780
tele-manas 1820
tempo 15
temporal lobe epilepsy 16
temporal-lobe pathology 94
temporal-lobe syndrome 1001
Tensilon (edrophonium) test 1233
Tension-type headache 348
terminal half-life 1780
tertiary amines 1355
tertiary prevention 1970
testing effect 1244
tetany 683
Tetracosactrin testing 654
Thallium SPECT 1094
Thase and Rush model 1485
The bedside cognitive examination 992
The circuit of Papez 967
the great imitator 275
the medial temporal lobe structures 1178
Therapeutic drug monitoring 257
therapeutic drug monitoring 257
therapeutic index 807
therapeutic threshold 1287
therapeutic window 1288
therapist effects 1715
theta rhythm 1603
theta-burst stimulation 1575
Thiamine 514
Thiamine pyrophosphate 713
thioxanthenes 658
three canonical vascular cognitive impairment subtypes 881
THREE-D 327
thrombotic limb 2088
thyroid-stimulating antibodies 671
to enhance human resource in mental health sub-specialties 1804
tonic firing of orexin neurons 381
tonic inhibition 1298
Topiramate 168
topiramate 168
topiramate titration 357
Topiramate-associated appetite suppression and weight loss 365
torsade de pointes 1564
total body surface area 2107
total tau 455
Tourette syndrome 1072
toxic epidermal necrolysis 246
TRAILBLAZER-ALZ 2 600
trained lay counsellor 1923
trait markers 329
transcranial magnetic stimulation 472
transdermal system 1423
transfer 91
Transformed chronic daily headache 350
Transient epileptic amnesia 981
transient global amnesia 981
transinstitutionalisation 1896
transition 77
transmission-risk behaviour 1089
transportability 1903
transporter preference 1349
transtentorial or foramen-magnum herniation 1035
trauma-induced alteration in mental status 892
Traumatic axonal injury 894

Traumatic cerebral contusions 893
treatment gap 49
treatment of schizophrenia in adults 1689
treatment-resistant depression 678
treatment-resistant schizophrenia 783
tremor entrainment test 1227
Treponema pallidum 1183
Treponemal tests 1184
tri-iodothyronine 1484
triage 533
tricyclic antidepressants 131
Trientine 701
triphasic waves 312
Triphasic waves 312
tRNS 1600
Trojan-horse 1081
trough 799
trough sample 1271
Tryptophan 728
tuberculoma 1190
tuberculous adrenalitis 652
tuberoinfundibular dopamine neurons 1295
tubular visual field 1227
twelve hour trough 1373
two-front 1241
Type B 194
Tyramine 1361
tyrosine hydroxylase 392

U

umbrella review 1705
unbound fraction 1382
uncompetitive open-channel NMDA receptor antagonist 1419
Uncoupling 388
undercount 1873
unintentional non-adherence 2043
unique users 1822
universal 172
unprovoked seizures 44
unresolved extraction 1678
unstable at meiosis 1136
uptake 354
uraemic encephalopathy 312
urine osmolality 771
utilisation 998
Utilisation behaviour 998

V

vagus nerve stimulation 121
valid denominator 1799
valproate 168
Vascular cognitive impairment 472
Vascular cognitive impairment without dementia 873
vascular dementia 315
vascular depression hypothesis 880
vasopressin V2 receptor antagonists 773
Vegetative state 294
ventral capsule/ventral striatum 1592
ventromedial frontal damage 977
vertical supranuclear gaze palsy 1215
vertical supranuclear ophthalmoplegia 707
VGKC-complex 1166
video-EEG telemetry 56
video-electroencephalography 24
vigabatrin 198
viral reservoir 1081
viscosity 171
Visual agnosia 999
visual hallucinations 424
Visual migraine aura 344
vitamin D insufficiency 685
VMAT2 inhibitors 1149
voltage-dependent sodium channels 1387
voltage-gated sodium channels 1381
voltage-sensitive sodium channels 1356
volume depletion 1372
Volume of distribution 1271
volume of distribution 1271
voluntary sector 1806

W

waist circumference 740
Wake-like cortical activity 376
washout 1023
water intoxication 773
water loading 782
Wernicke encephalopathy 514
Wernicke's encephalopathy 711
what a service has 1813
what a service owes 1813
Whipple's disease 1196
white matter dementia 827
wide therapeutic index 1286
Wilson disease 514
Wilson protein 688
Wilson's disease 631
wing beating 693
within-class selection 1433
Witzelsucht 987
working memory 449
written confirmation 1660

X

xanthochromia 320

Y

young-onset dementia 426

Z

Zero-order kinetics 1272
Zinc acetate 701
zona fasciculata 636
zona glomerulosa 636
zonisamide 168

#

“migraine personality” 344

About the authors

Dr. Wilfred D'souza

MD Psychiatry (Resident) · Certified Gestalt Therapist · Weave, Centre for Integrative Psychiatry

Works in integrative psychiatry with Schema Therapy as the primary lens, alongside CBT, DBT, motivational interviewing and emotion-focused therapy, and is a certified Gestalt therapist. His clinical interests run through neurodivergence, ADHD, bipolar disorder, personality disorders, addiction and developmental psychiatry: work grounded in lived experience of attention difficulties long labelled as laziness, which made psychiatry personal before it became professional. He has been developing the Toroidal Model of Psychosis and Affect (TMoP), his computational model of mind, and TMoP thinking shaped the Claude Code AI system he builds with: clinical software, teaching tools and much of this series' production pipeline among its work, bringing psychiatry into the AI era in his own way. Lead author of the Weave Sprint series.

Dr. Niharika Reddy

Consultant Psychiatrist · Weave, Centre for Integrative Psychiatry

Consultant Psychiatrist at Weave, Mumbai, seeing people online through weave.clinic. Her work centres on relationships and family dynamics, family therapy and women's mental health, with a particular focus on neurodivergence in women and on adolescent mental health. Practises integrative, person-centred psychiatry. Co-author of this volume.

THE WEAVE LIBRARY

You revised from one. The whole shelf is open.

Field guides and references written in practice, for the people who do the work and the people living through it. Most are free to read and download.



field guides · references · for clinicians and families · free to read

weave.clinic/library

TAP OR SCAN TO OPEN

WEAVE STUDY GUIDES

When the night is short, read the guide.

Twenty-two topic guides across all four papers. Study notes, model answers, mnemonics, high-yield comparisons and previous-question frequency. Free, searchable, built for this exam.



22 guides · all four papers · fully searchable · free to read

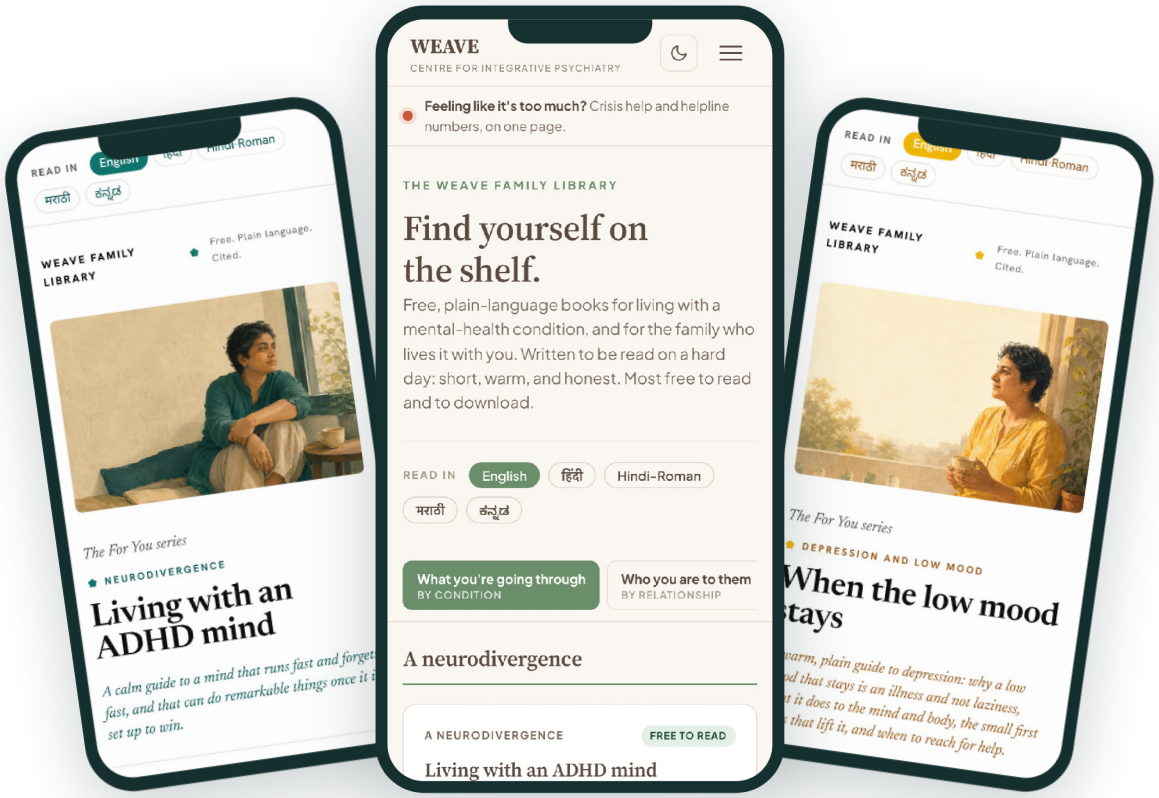
weave.clinic/study

TAP OR SCAN TO OPEN

THE WEAVE FAMILY LIBRARY

Send the family home with more than a prescription.

Free, plain-language books for living with a mental-health condition, written for patients and the people who love them. Find yours by condition, by relationship, or by situation.



plain language · patients and families · five languages · free to read

weave.clinic/family

TAP OR SCAN TO OPEN

A PAPER A DAY

The whole written exam, in one day's reach.

Every paper of the MD psychiatry written exam in one free book, with a colour memory code and clickable navigation, for the last week, the last night, the last hour before the hall.



four papers · one volume · colour-coded · free to read

weave.clinic/revision

TAP OR SCAN TO OPEN

STUDY TOGETHER

You do not have to revise alone.

These notes started as messages in a room full of residents helping each other through. That room is still open. Bring your doubts, your corrections, and the questions the textbooks do not answer.

Join the study group

https://chat.whatsapp.com/H7TsLGK2wFl4bSe2r0Zfzn?mode=gi_t

PRINTING AND DISTRIBUTION

Notes Next Door (NND), Delhi is the official printing and distribution partner for Weave Sprint. The soft-copy PDFs and the complete print-ready set are always available at weave.clinic/sprint, should you prefer to print the notes yourself. Weave Sprint is open source and free, and as part of Weave's commitment to residents it will stay that way, forever. We do, however, strongly recommend a printed set from Notes Next Door: in our experience, they have been the most reliable printing partner we have worked with.

WEAVE SPRINT · PAPER IV

The fourth paper of the MD psychiatry written examination, gathered into one book: the neurological interface from epilepsy and the dementias through stroke, injury and delirium; the endocrine, metabolic and systemic diseases that present to psychiatry before they present anywhere else; psychopharmacology in full, from why a drug is chosen to what to do when it fails; the recent advances and the studies behind them; and psychiatry at the scale of a population. Eleven complete sets of notes, written to be understood the first time and recalled every time after, with a colour memory code, a clickable contents and a full index.

Built to understand. Kept to revise.

All the very best for your exam.

WEAVE

weave.clinic · wilfred@weave.clinic · niharika@weave.clinic