WEAVE SPRINT
PAPER IV · VOLUME FIVE

MD and DNB psychiatry, theory revision

Neurology, Medicine and Recent Advances

Psychopharmacology, Principles and Emergencies

How the drugs work and what that predicts, then what to do when they fail or harm: the emergencies, resistance, monitoring and interactions.

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, Volume Five: Psychopharmacology, Principles and Emergencies

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 fifth volume is psychopharmacology given two chapters, because the second one is where most of the difficulty lives. The first sets out how the drugs work and what that lets you predict before you prescribe. The second is everything after: the emergency, the patient who does not respond, the monitoring that should have caught it, and the interaction that caused it.

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 VII **Psychopharmacology I: Principles and Core Drug Classes** 11

- Receptor pharmacology and pharmacokinetic principles 11
- Antipsychotics 47
- Antidepressants 93
- Mood stabilisers 126
- Anxiolytics, hypnotics and other agents 158
- Consolidate and apply 186
- Sources and further reading 198

PART VIII **Psychopharmacology II: Emergencies, Resistance, Monitoring and Interactions** 201

- Drug emergencies 201
- Treatment resistance and augmentation 237
- Drug monitoring, pharmacovigilance and pharmacogenomics 267
- Interactions and prescribing in physical illness 290
- Neurostimulation and device-based treatments 329
- Pharmacogenetic and structural drug-reaction syndromes 368
- Consolidate and apply 383

BACK MATTER

Subject index 405

About the authors 410

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

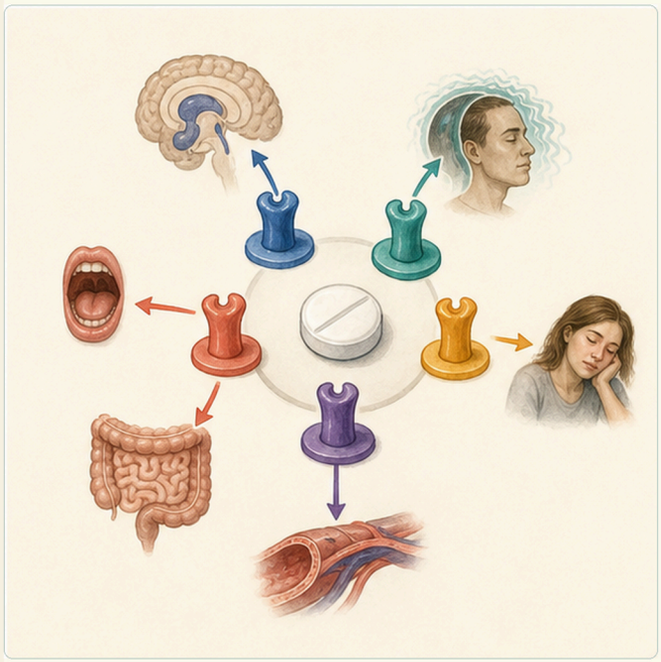
05 Anxiolytics, hypnotics and other agents

06 Consolidate and apply

07 Sources and further reading

PART 7 • Psychopharmacology I: Principles and Core Drug Classes

VISUAL 7.1



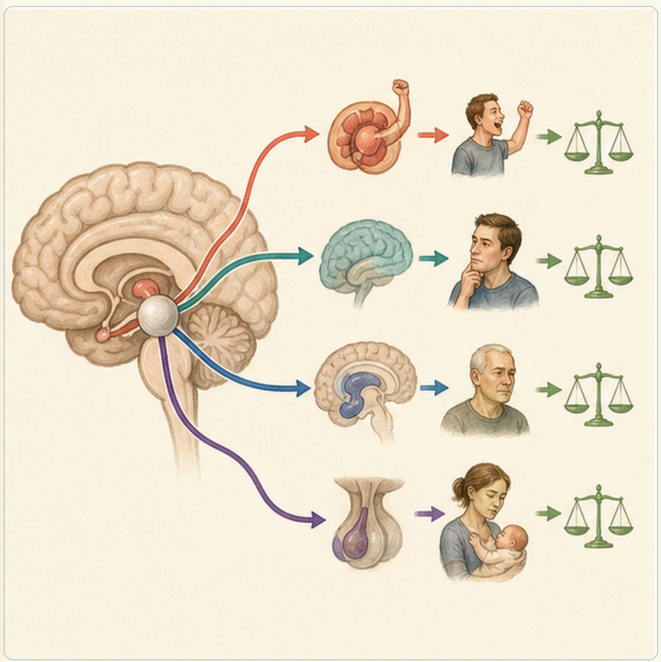
A receptor profile predicts a burden profile

- 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



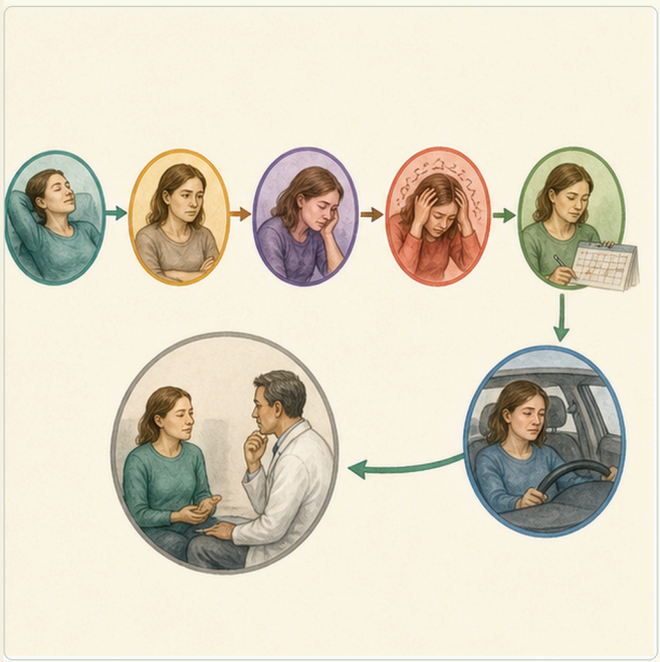
Lithium safety is a monitoring system

- 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.

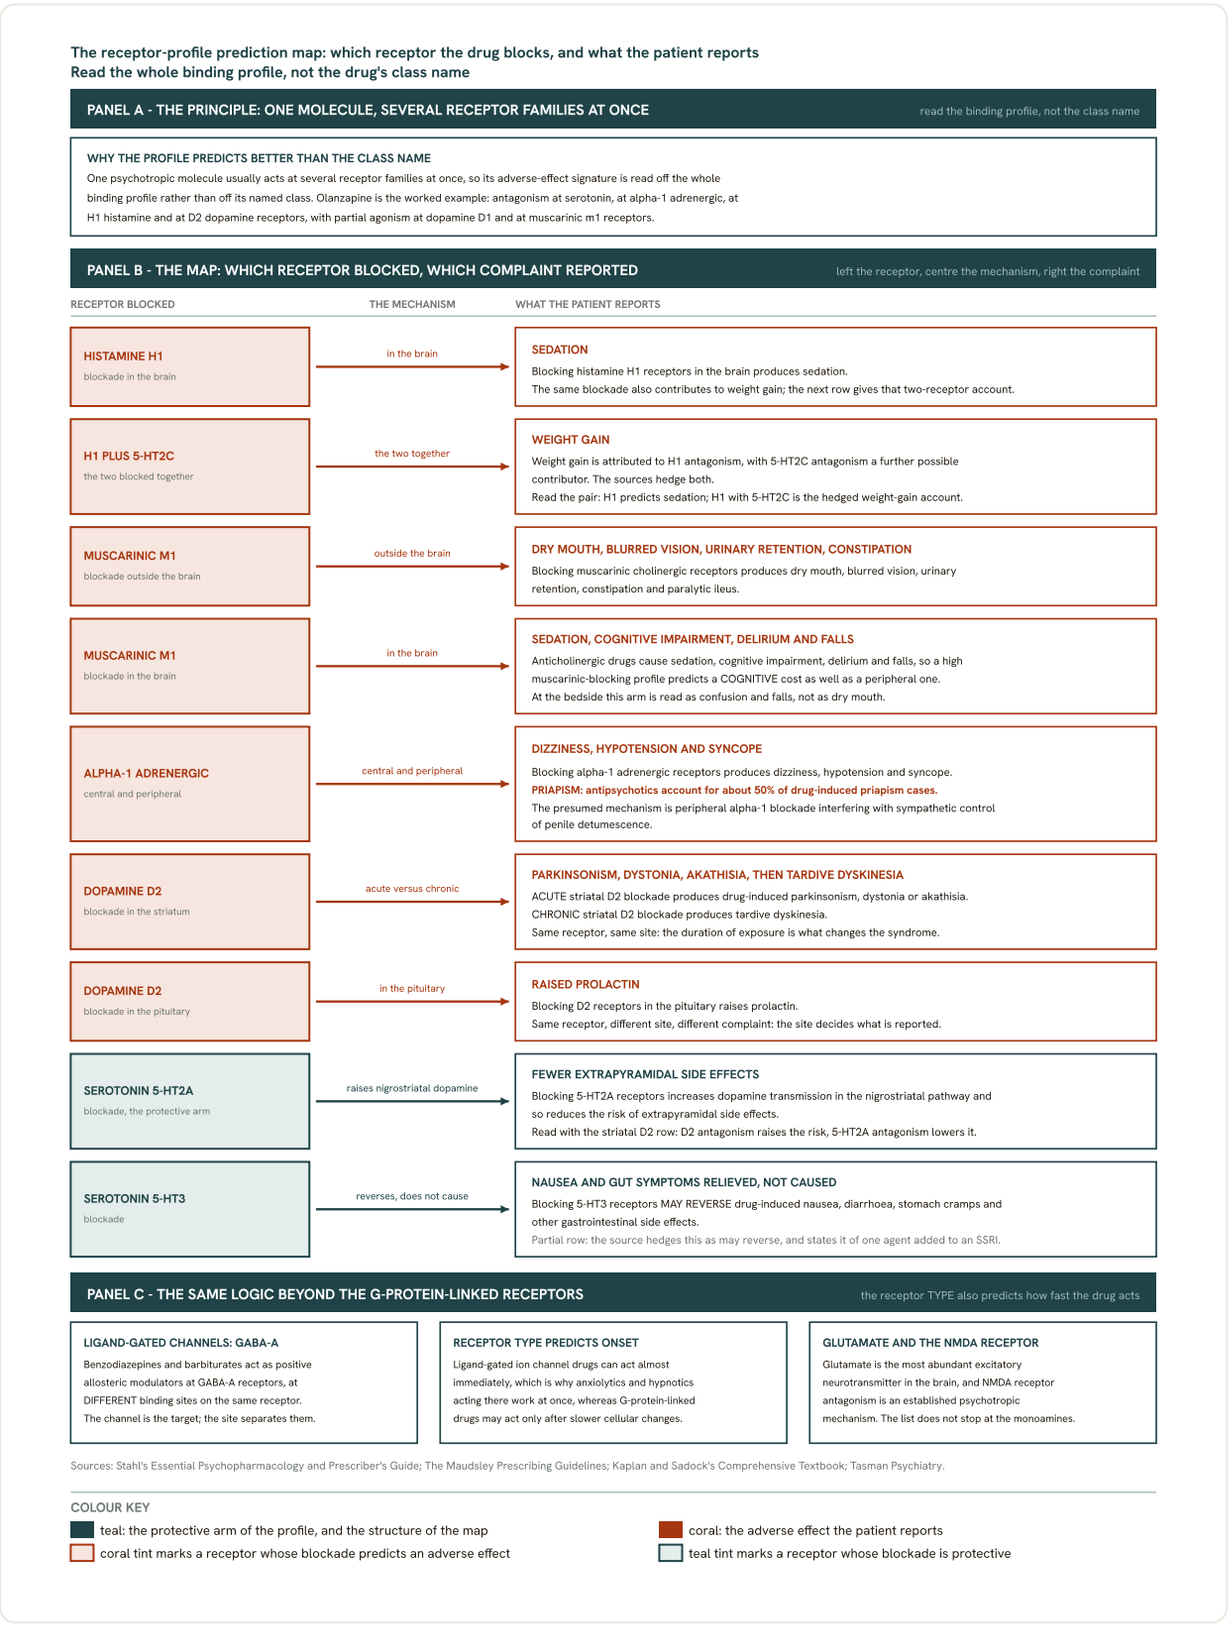


Fig 36.1 **The receptor-profile prediction map.** The chapter's organising figure: the adverse effect a patient reports is read off the drug's whole binding profile, not off its class name. H1 blockade predicts sedation, H1 with 5-HT2C weight gain, muscarinic blockade a peripheral and a cognitive cost, 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 LOWEST THERAPEUTIC INDICES AMONG DRUGS IN USE	**More than 100** THE SAFE END OF THE SAME SCALE	**25 percent** RESPOND BELOW THE PUBLISHED TARGET RANGE	**Up to 25 percent** 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.

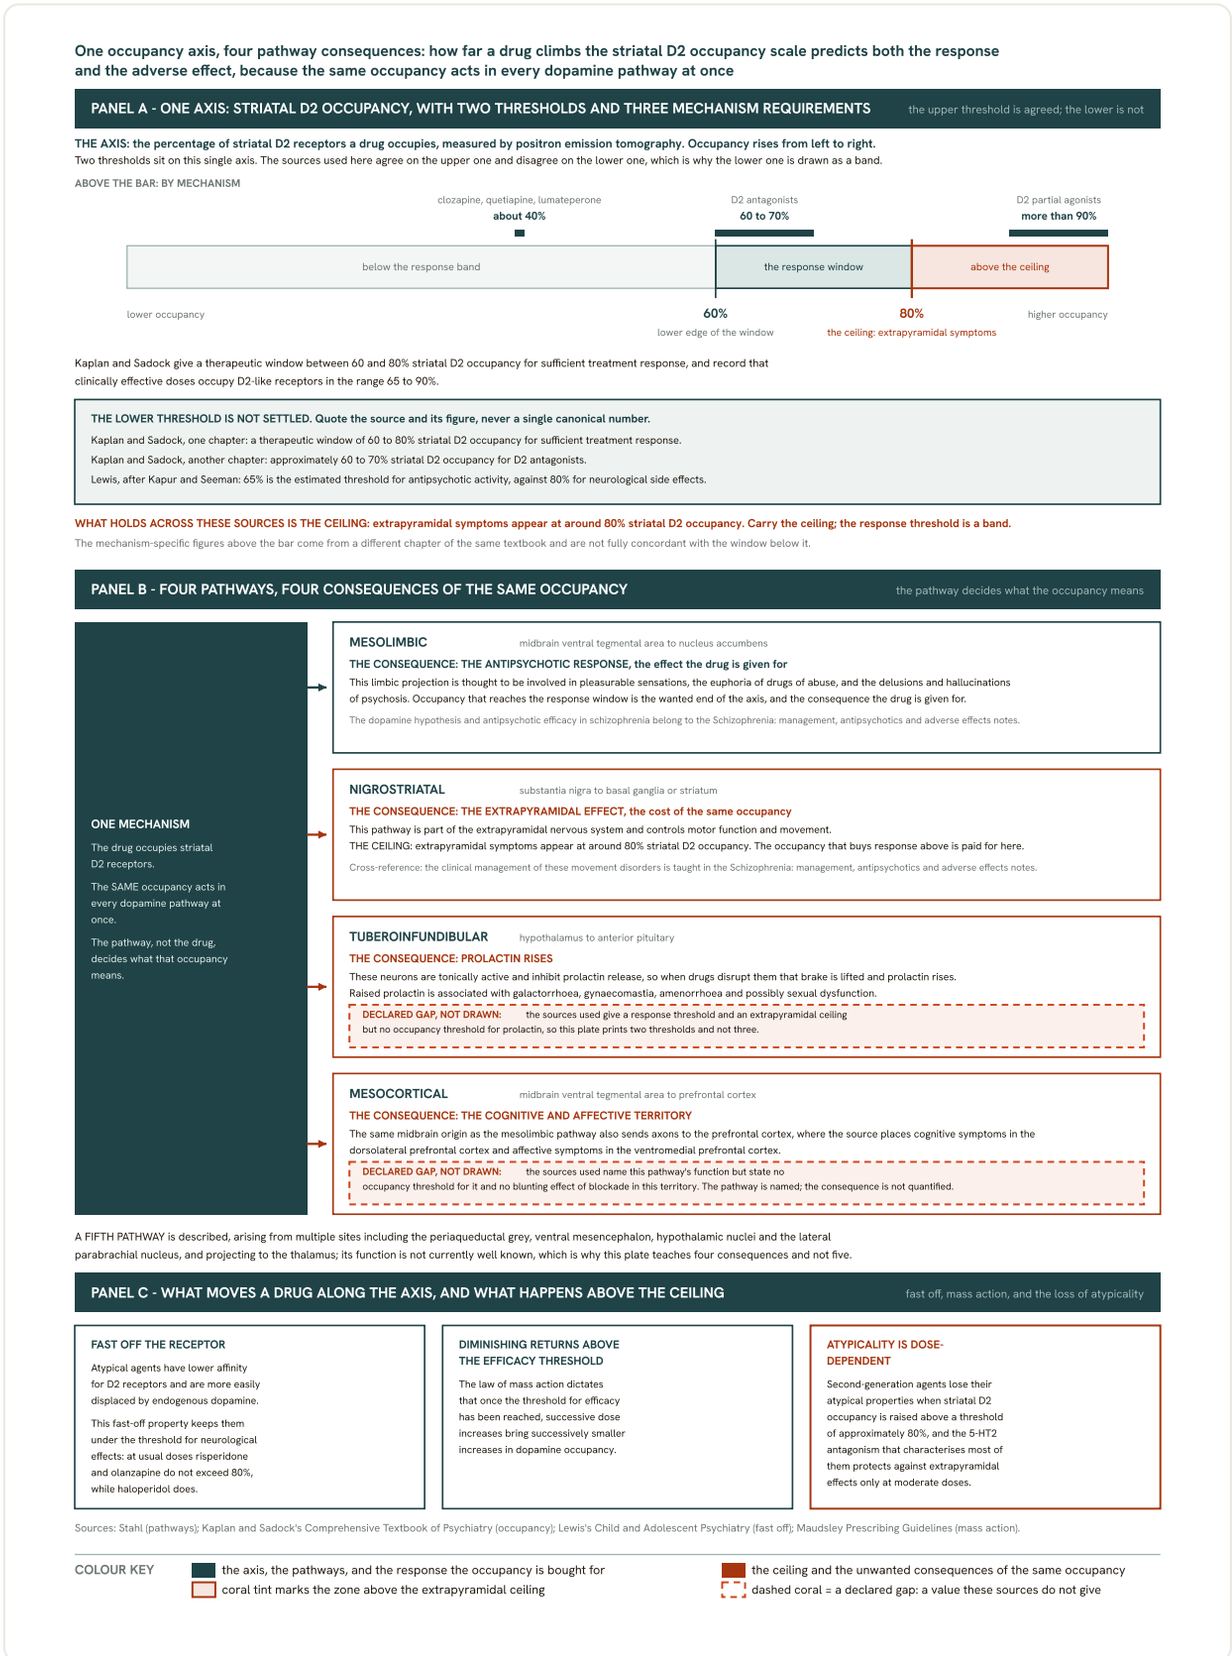


Fig 36.2 **One occupancy axis, four pathway consequences.** How far a drug climbs the striatal D2 occupancy scale predicts both its response and its adverse effects: the sources here agree on a ceiling of around 80% above which extrapyramidal symptoms appear, and disagree on the response threshold, placing 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.

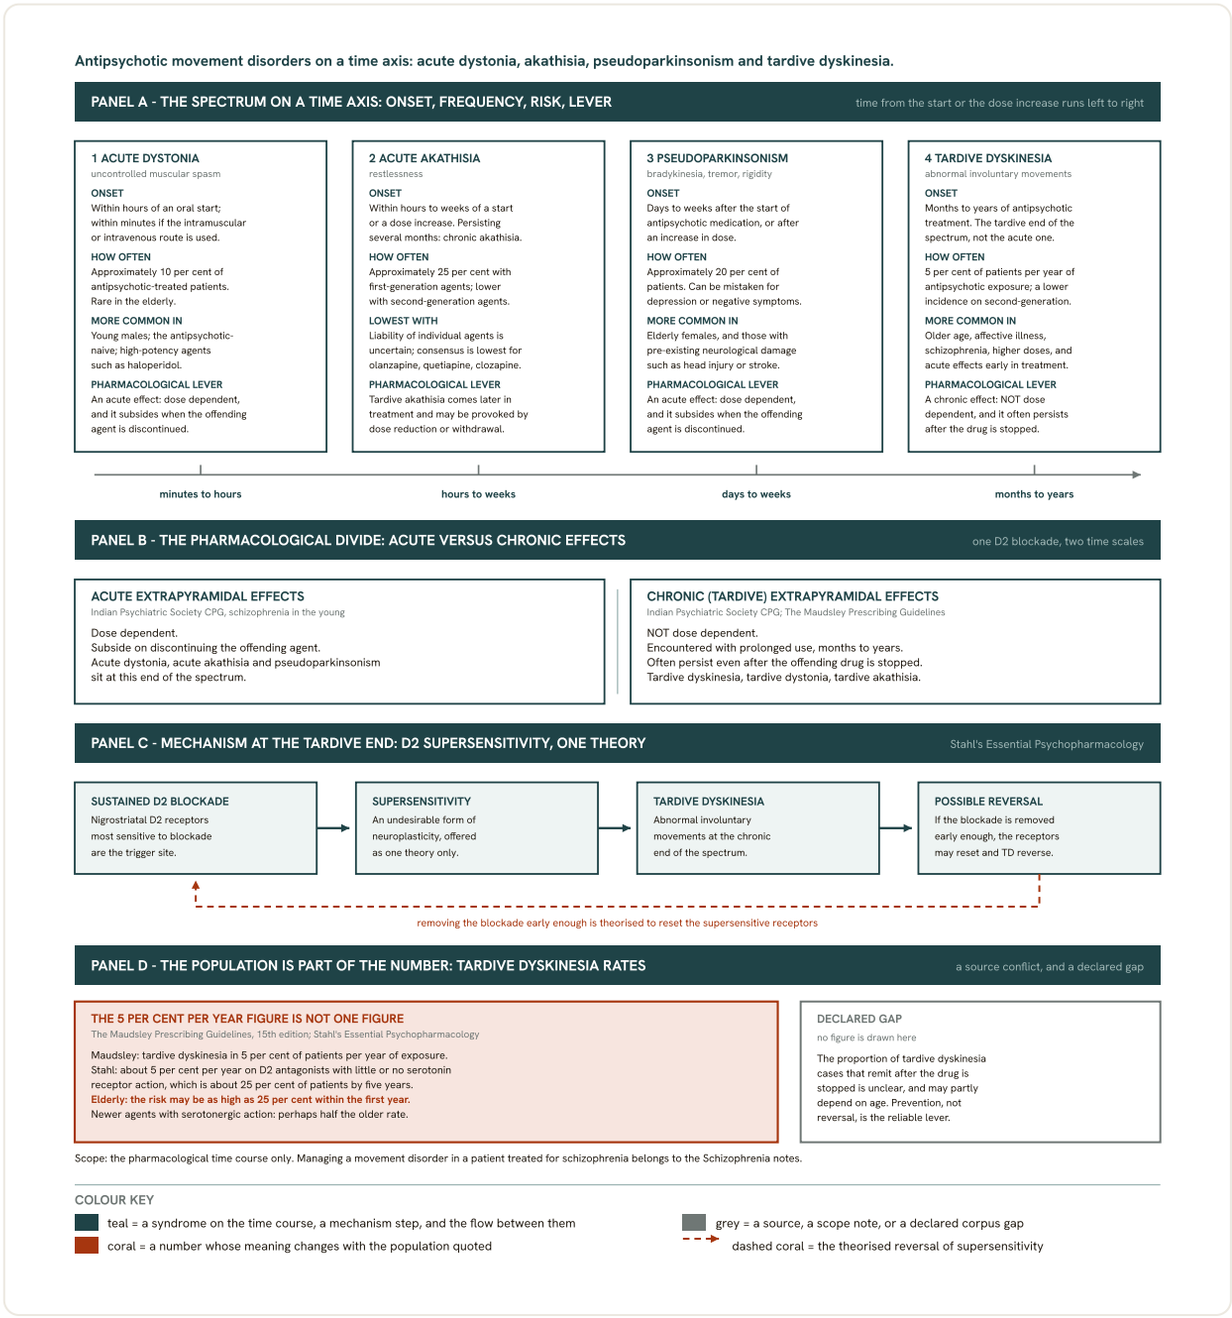


Fig 36.3 **The antipsychotic movement-disorder spectrum read on a time axis.** The four syndromes separate first by when they appear: acute dystonia within hours of an oral start and within minutes by the intramuscular or intravenous route (approximately 10 per cent of patients, more common in young 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.

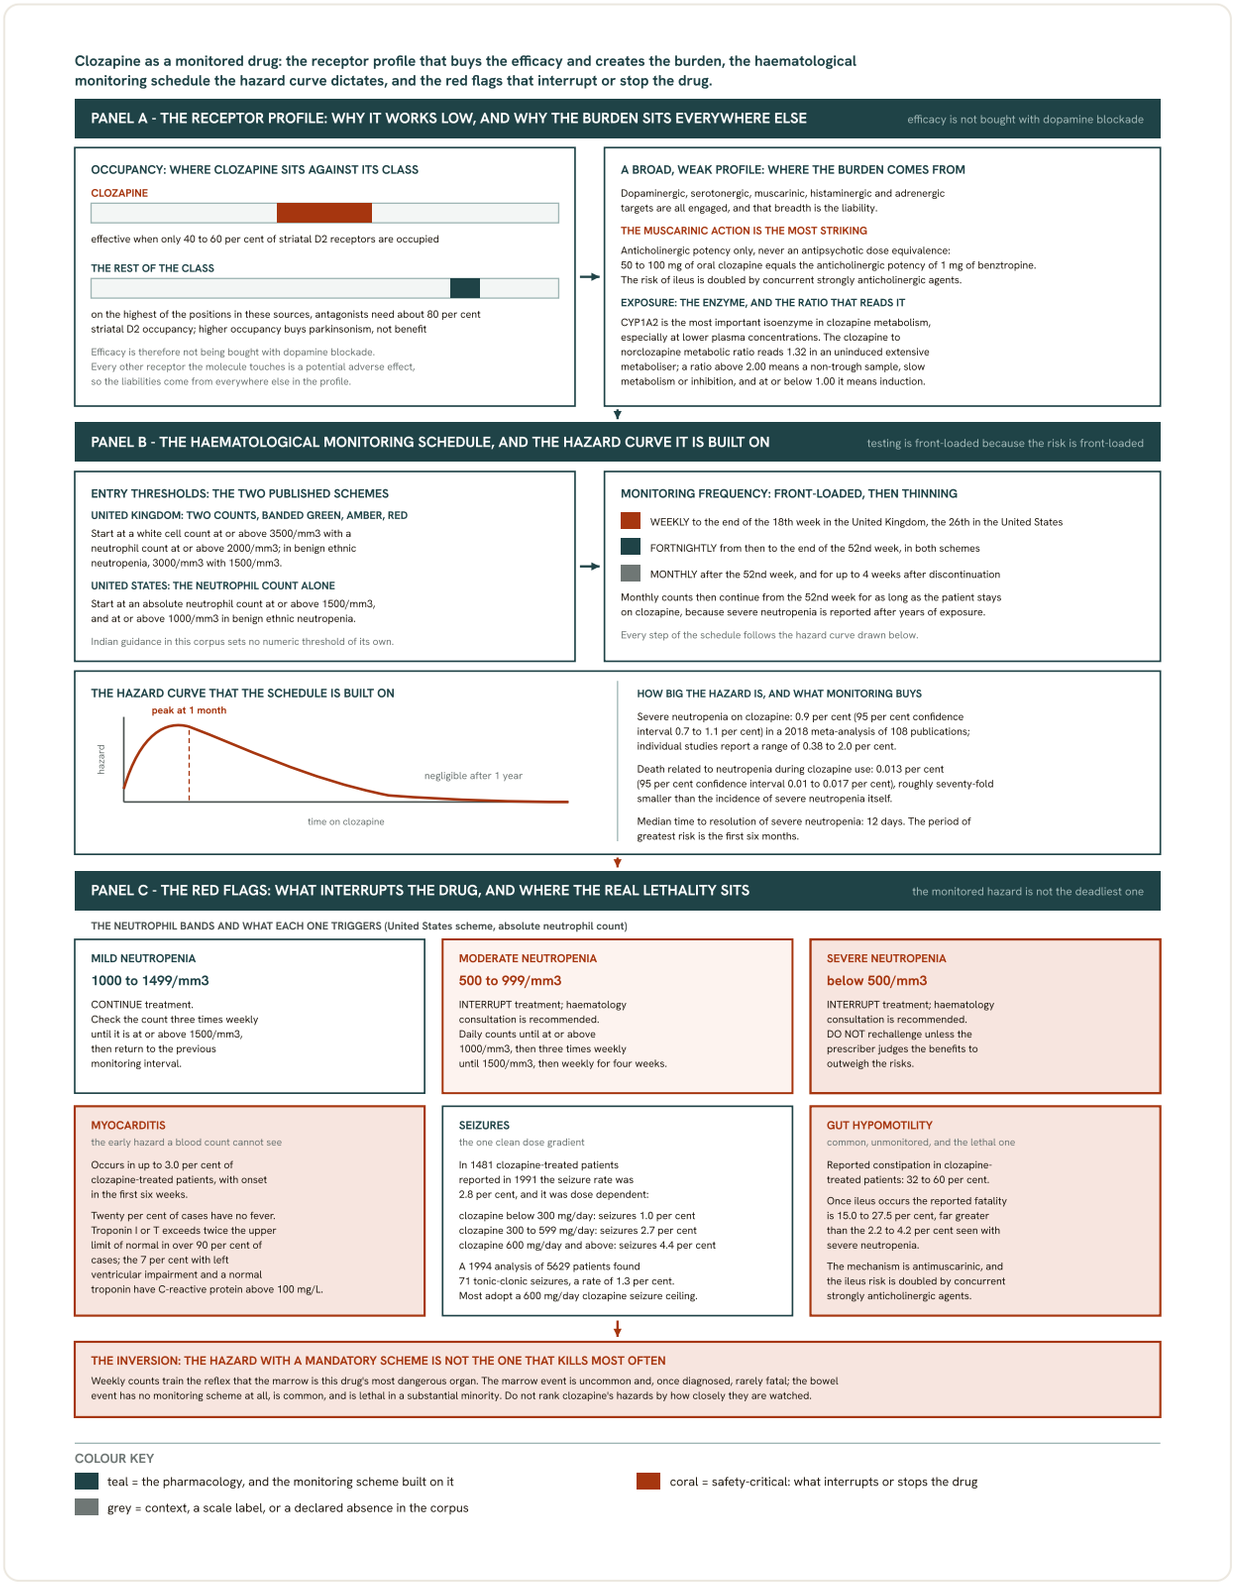


Fig 36.4 Clozapine: the profile that buys the efficacy, the monitoring schedule the hazard curve dictates, and the red flags that stop the drug. Clozapine is effective when only 40 to 60 per cent of striatal D2 receptors are occupied, where the antagonists as a group need approximately 80 per cent, so 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

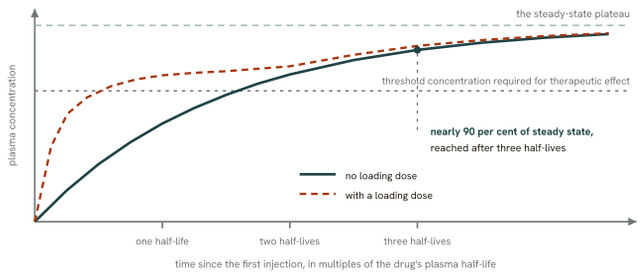
Nothing about a depot's opening months behaves the way a prescriber's instinct expects. With most long-acting injectable preparations, plasma concentrations climb for several weeks to months with no increase in dose at all, because each injection is given before the previous one has been cleared. That 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.

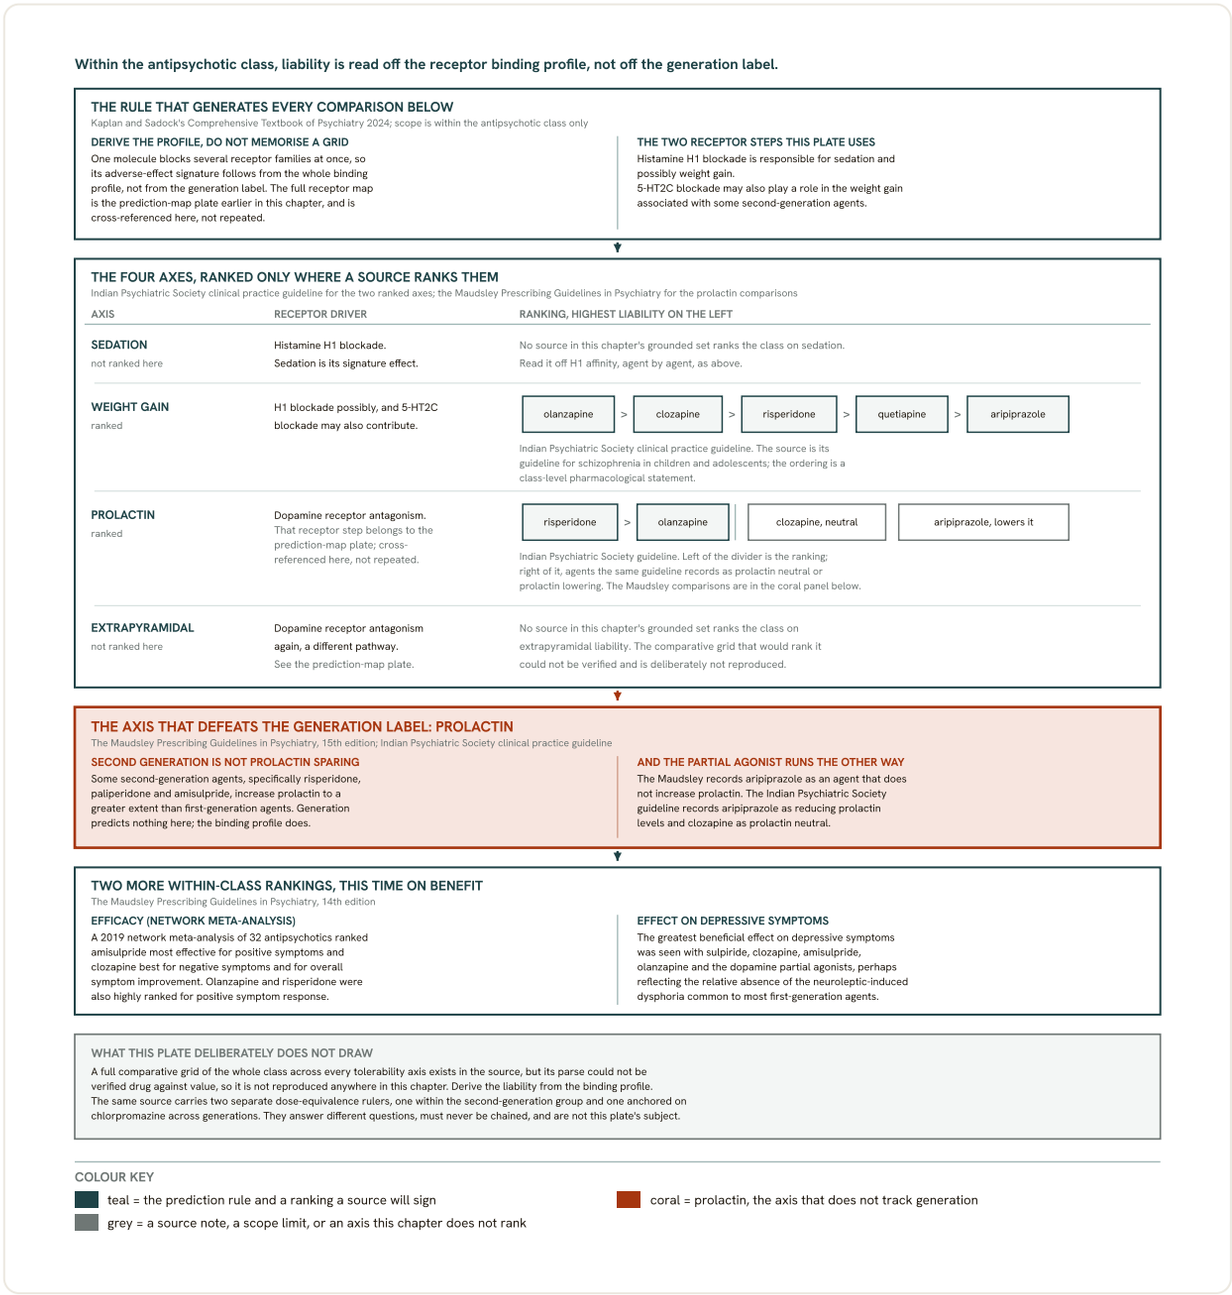


Fig 36.6 Within the antipsychotic class: what the receptor profile predicts, and where a source will actually rank the agents. Weight gain and prolactin are ranked where the Indian Psychiatric Society guideline ranks them, while sedation and extrapyramidal liability are drawn as unranked, because no 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.

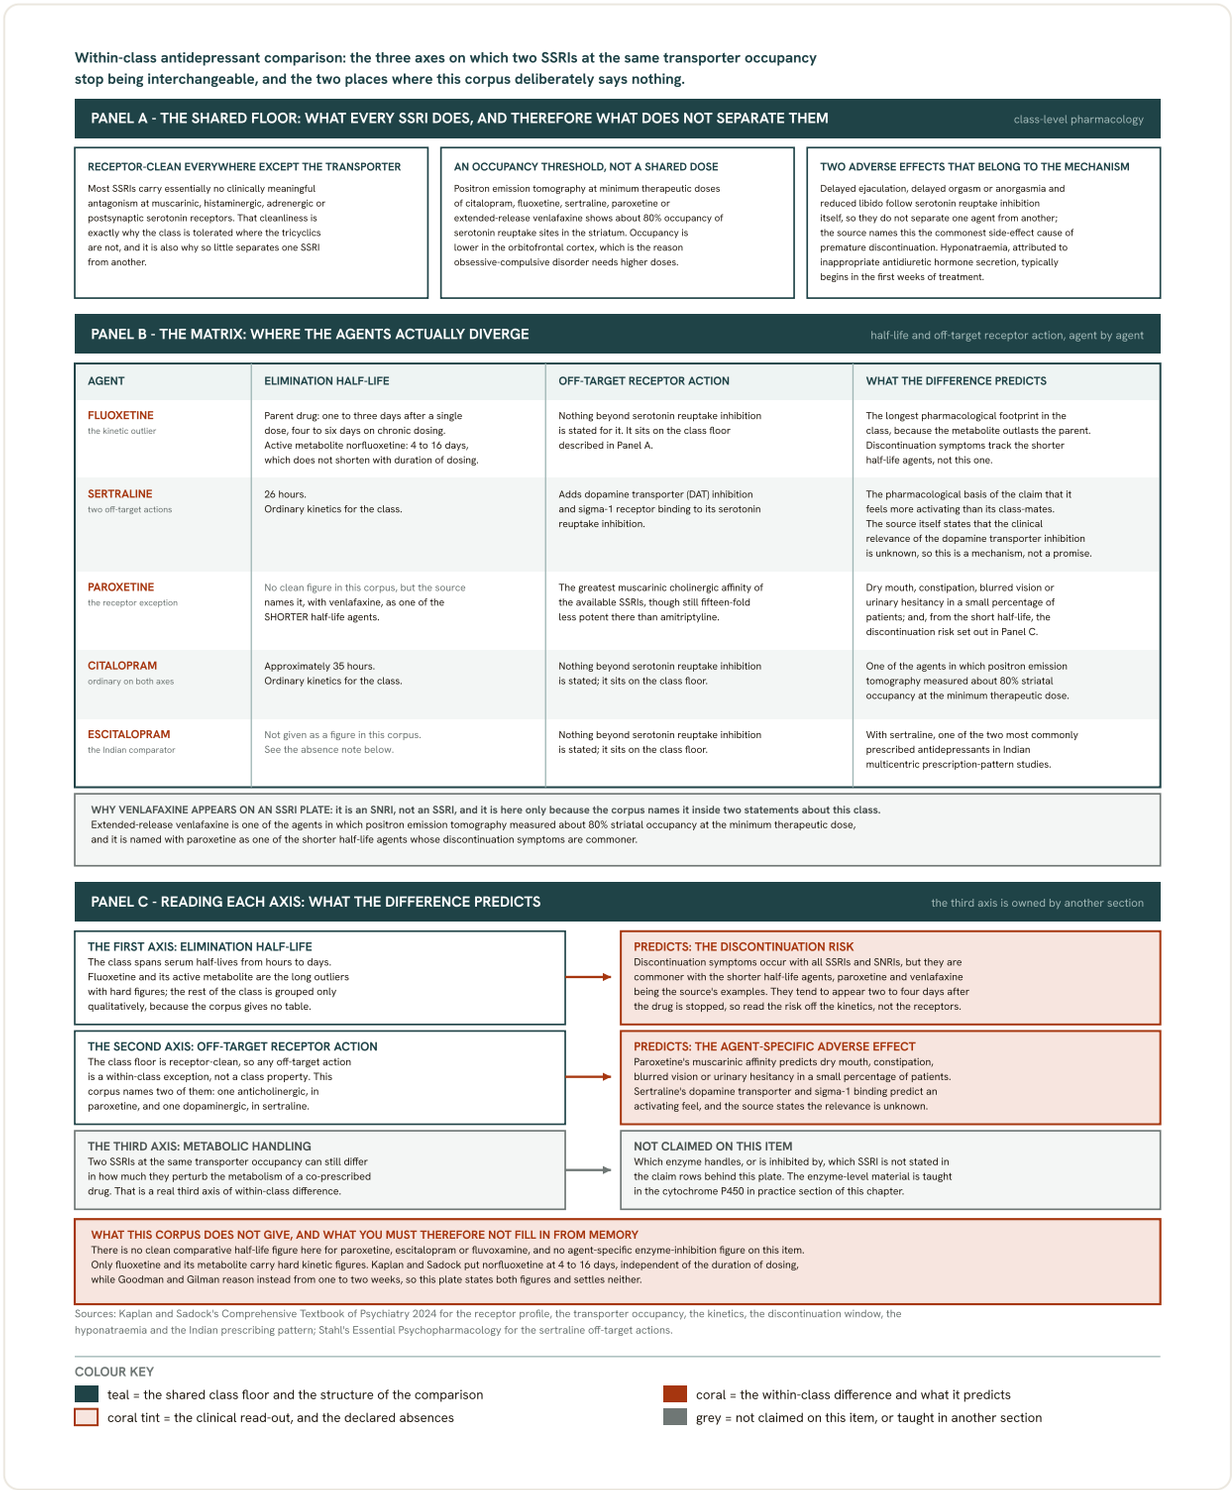


Fig 36.7 Within-class antidepressant comparison: half-life, off-target receptor action, and the axis this item does not claim. Most SSRIs are receptor-clean beyond the transporter, and every agent has to reach an occupancy threshold rather than a shared milligram dose, which is why sexual dysfunction 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 PERIPHERAL HALF-LIFE AT STEADY STATE	**28 to 48 hours** HALF-LIFE IN THE BRAIN	**80 to 100 percent** BIOAVAILABILITY, ANY PREPARATION	**10 to 14 hours** 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.

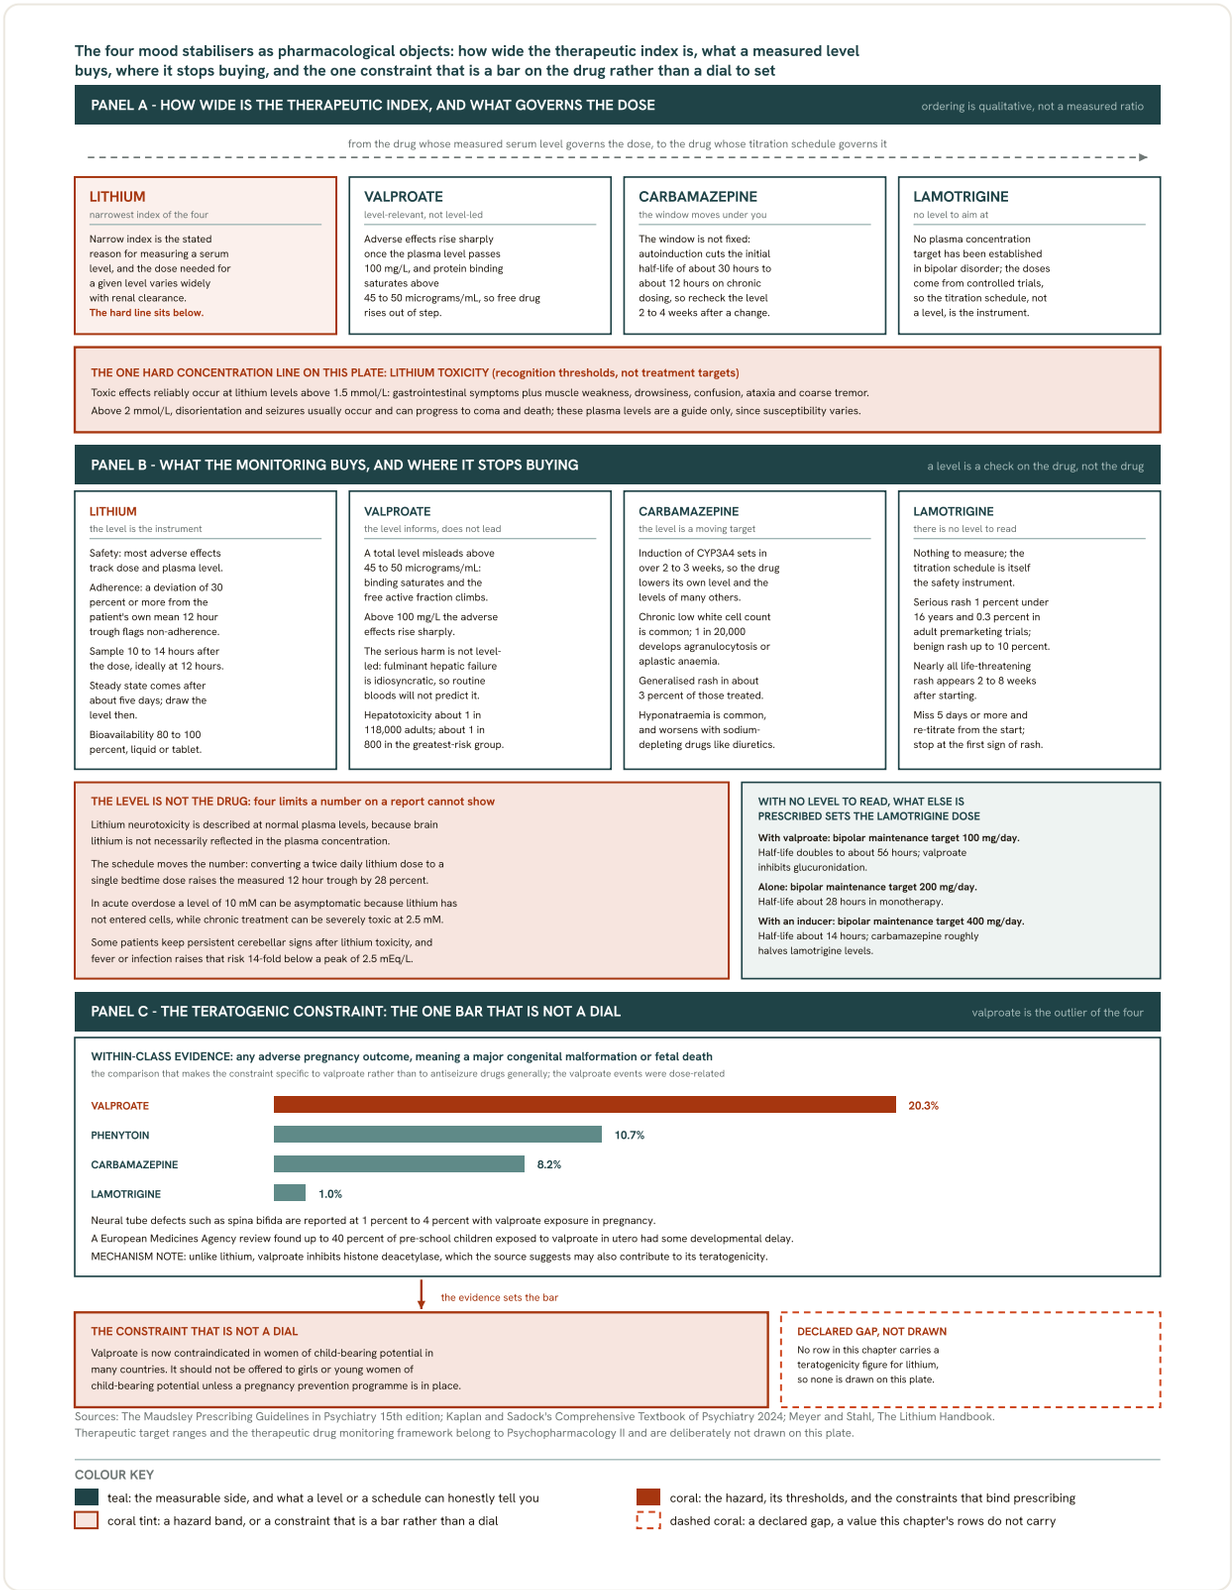


Fig 36.8 **The four mood stabilisers compared as pharmacological objects.** Lithium sits at the narrow end, where the serum level governs the dose and toxicity is reliable above 1.5 mmol/L; lamotrigine sits at the wide end, where no plasma target exists and the titration schedule is itself the safety 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	above 2 mmol/L	sevenfold	14-fold
TOXIC EFFECTS RELIABLY APPEAR	SEIZURES; CAN PROGRESS TO COMA AND DEATH	HOSPITALISATION RISK, ACE INHIBITOR IN THE ELDERLY	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 TIME TO PEAK PLASMA CONCENTRATION	**10 to 16 hours** PLASMA HALF-LIFE IN ADULTS	**20.3 percent** ANY ADVERSE PREGNANCY OUTCOME	**up to 40 percent** 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.

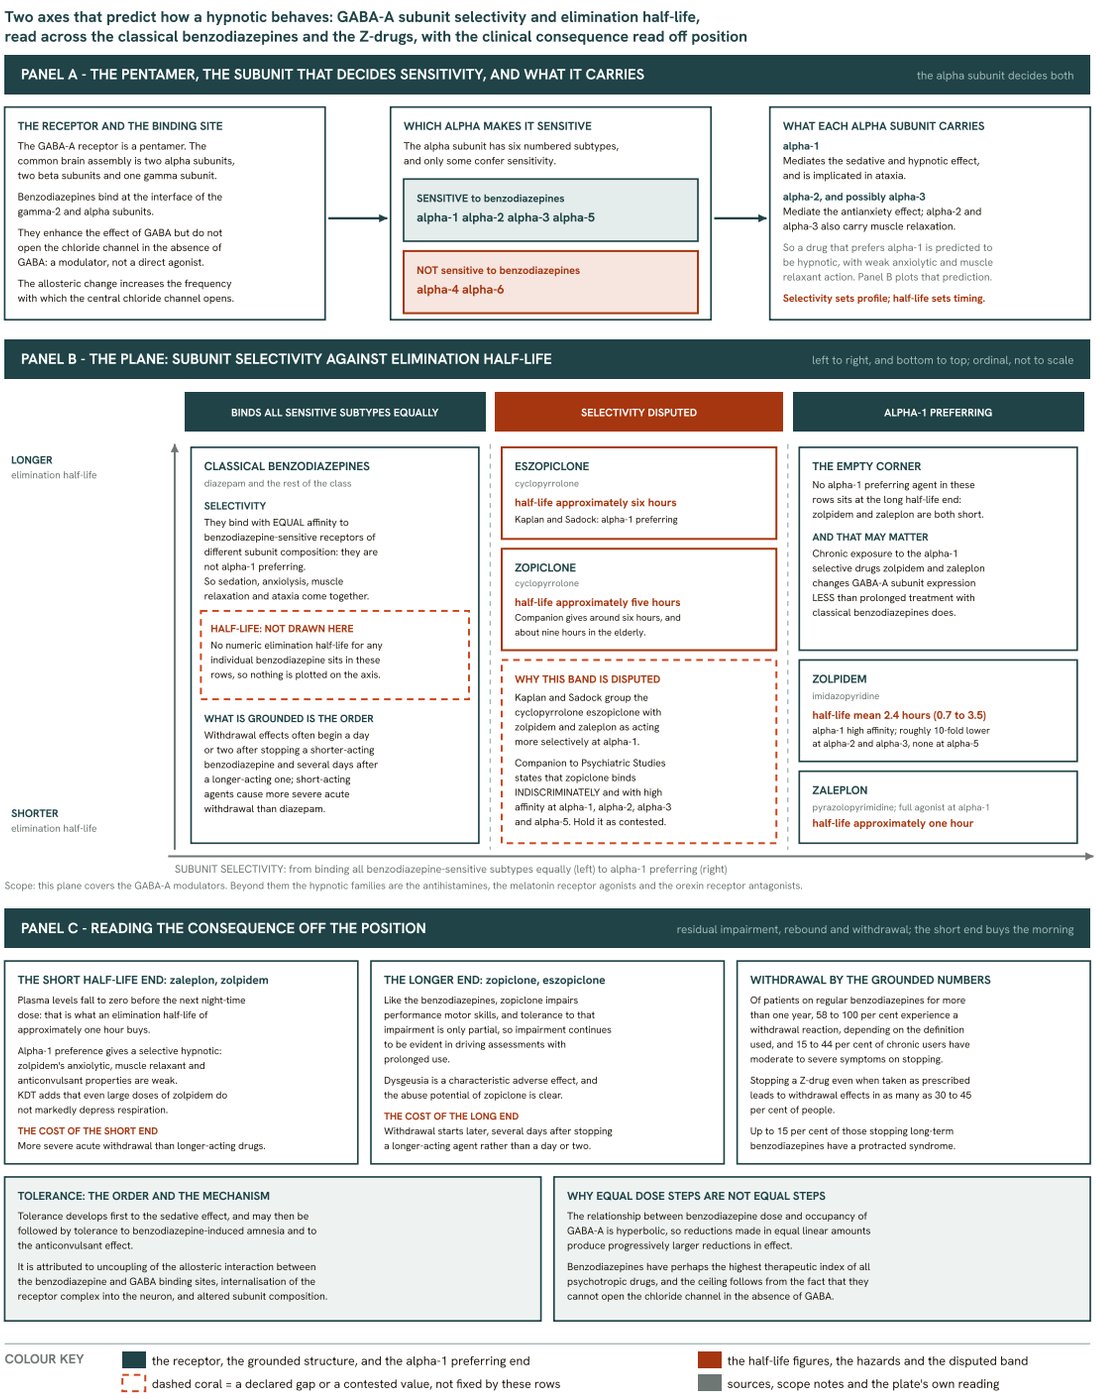


Fig 36.9 GABA-A subunit selectivity set against elimination half-life. Subunit preference predicts the profile and half-life predicts the timing: alpha-1 carries sedation, hypnosis and ataxia, while alpha-2 and possibly alpha-3 carry anxiolysis and muscle relaxation. The classical benzodiazepines bind 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

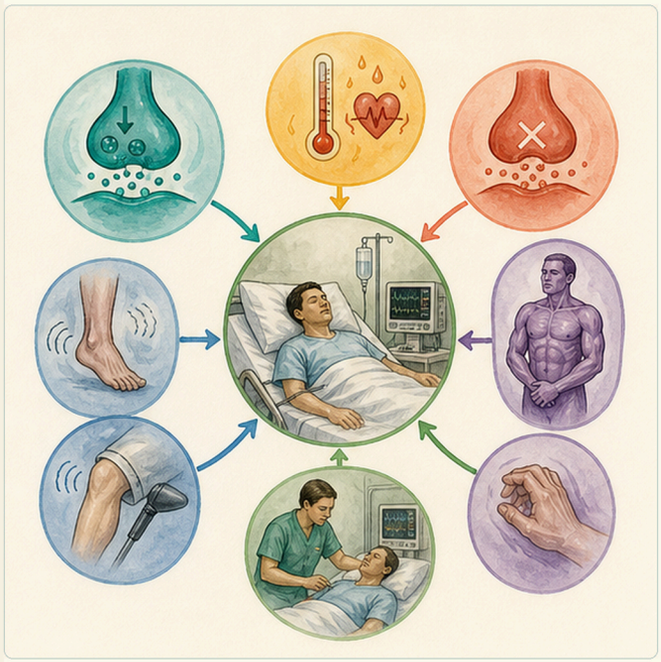
04 Interactions and prescribing in physical illness

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

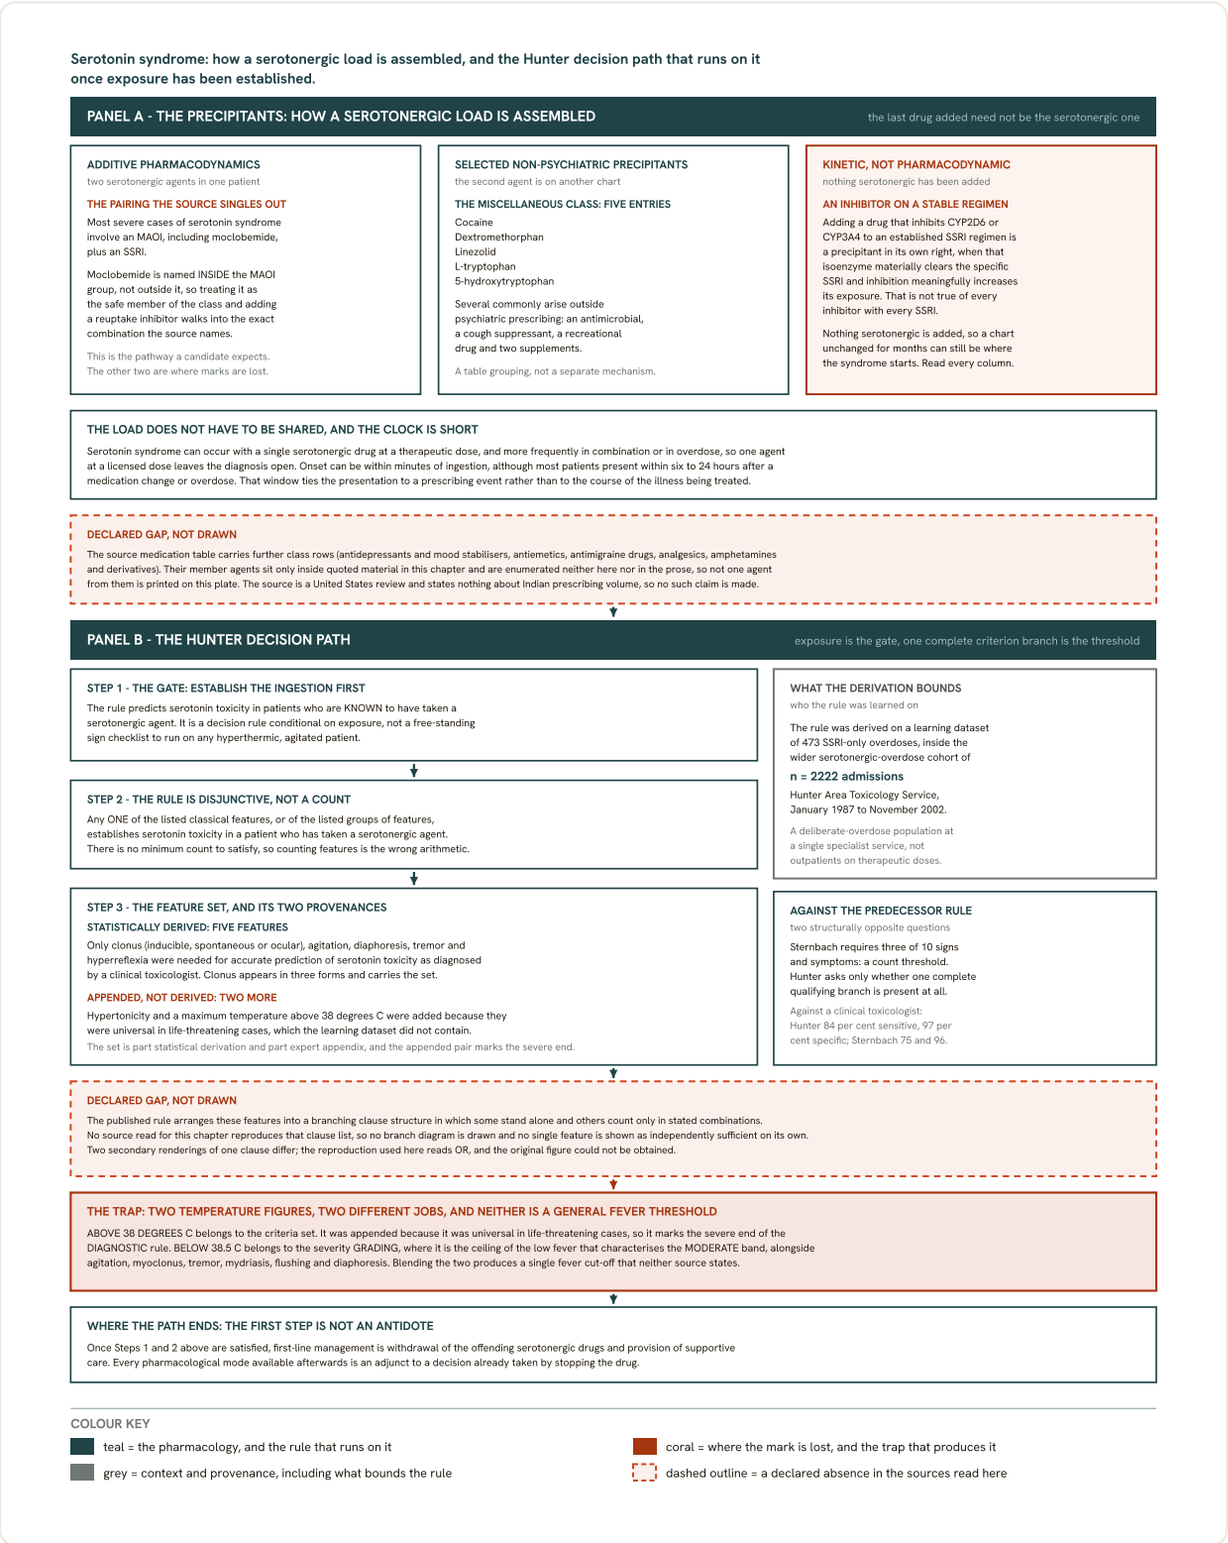


Fig 37.1 Serotonin syndrome: how a serotonergic load is assembled, and the Hunter decision path that runs on it. A single agent at a therapeutic dose leaves the diagnosis open, and most patients present within six to 24 hours of a medication change or overdose. Most severe cases involve an MAOI, including 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.

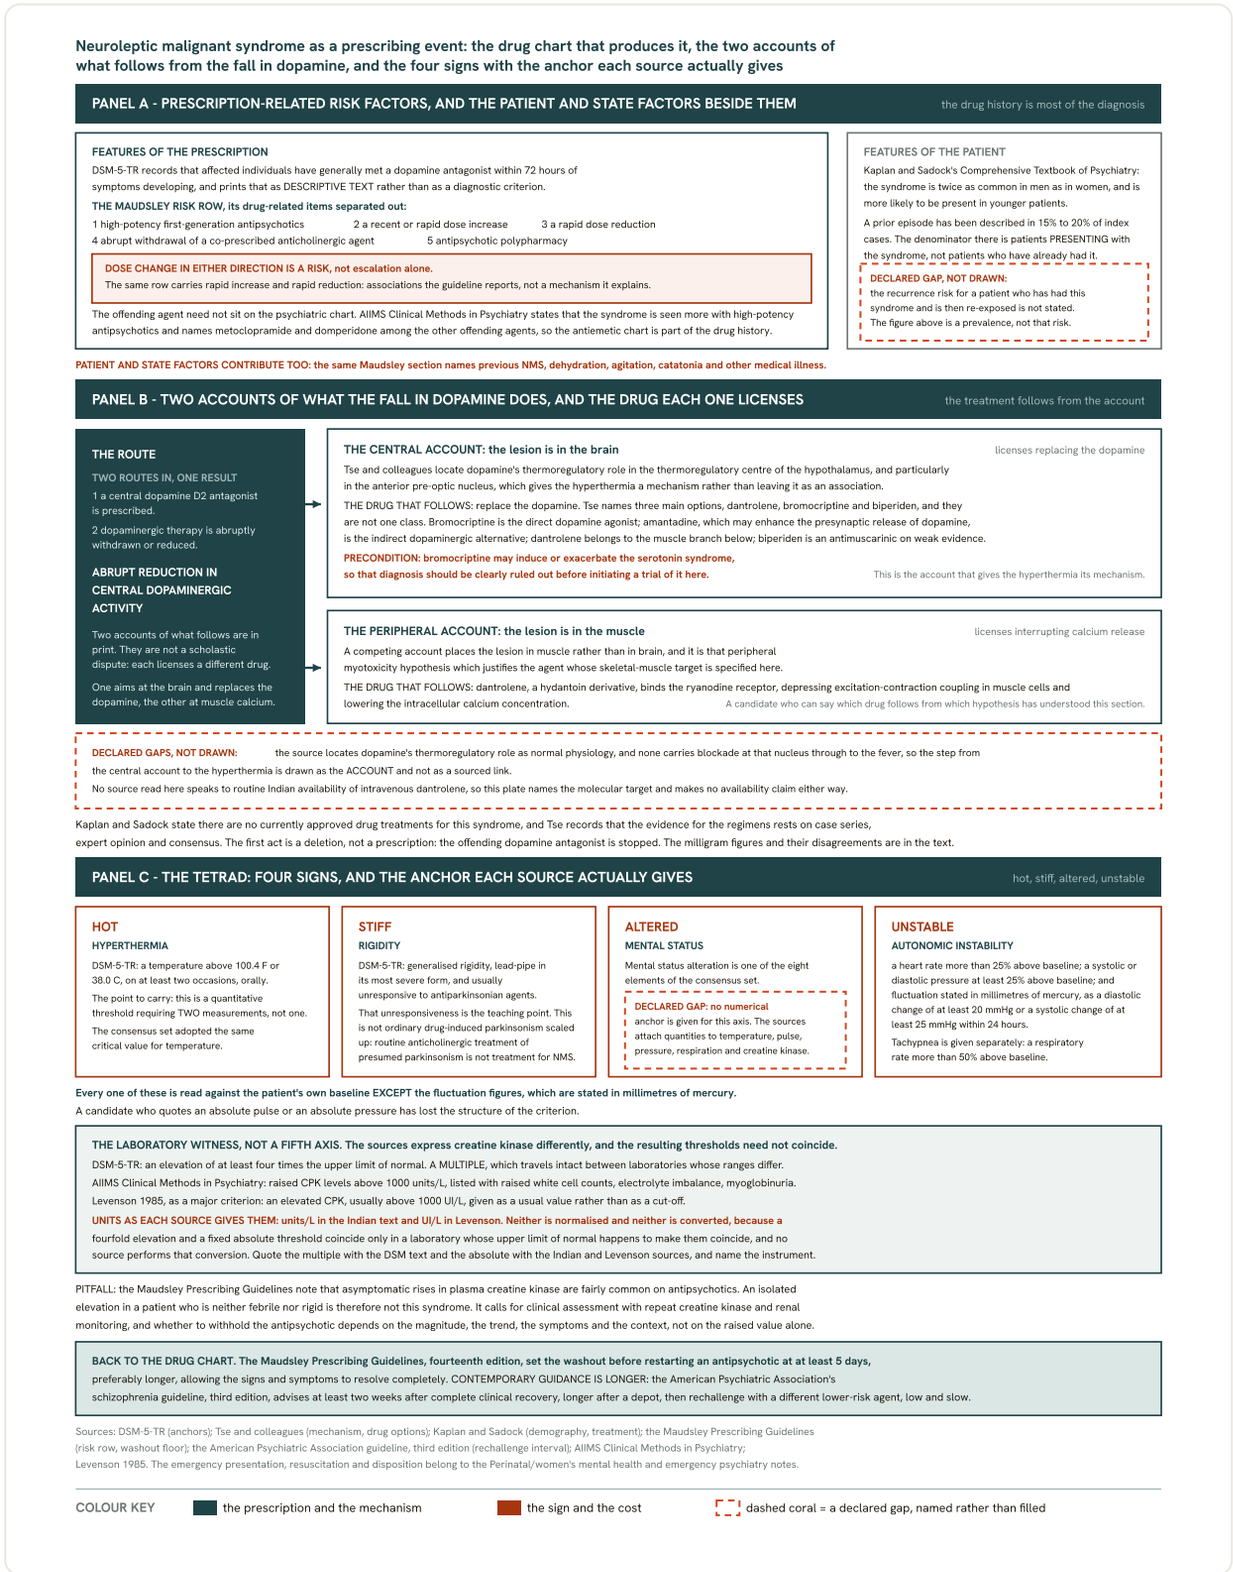


Fig 37.2 Neuroleptic malignant syndrome: the prescription, the two accounts, and the tetrad. Prescription-related risk factors: high-potency first-generation antipsychotics, a rapid dose change in either direction, abrupt withdrawal of a co-prescribed anticholinergic and antipsychotic polypharmacy, with 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

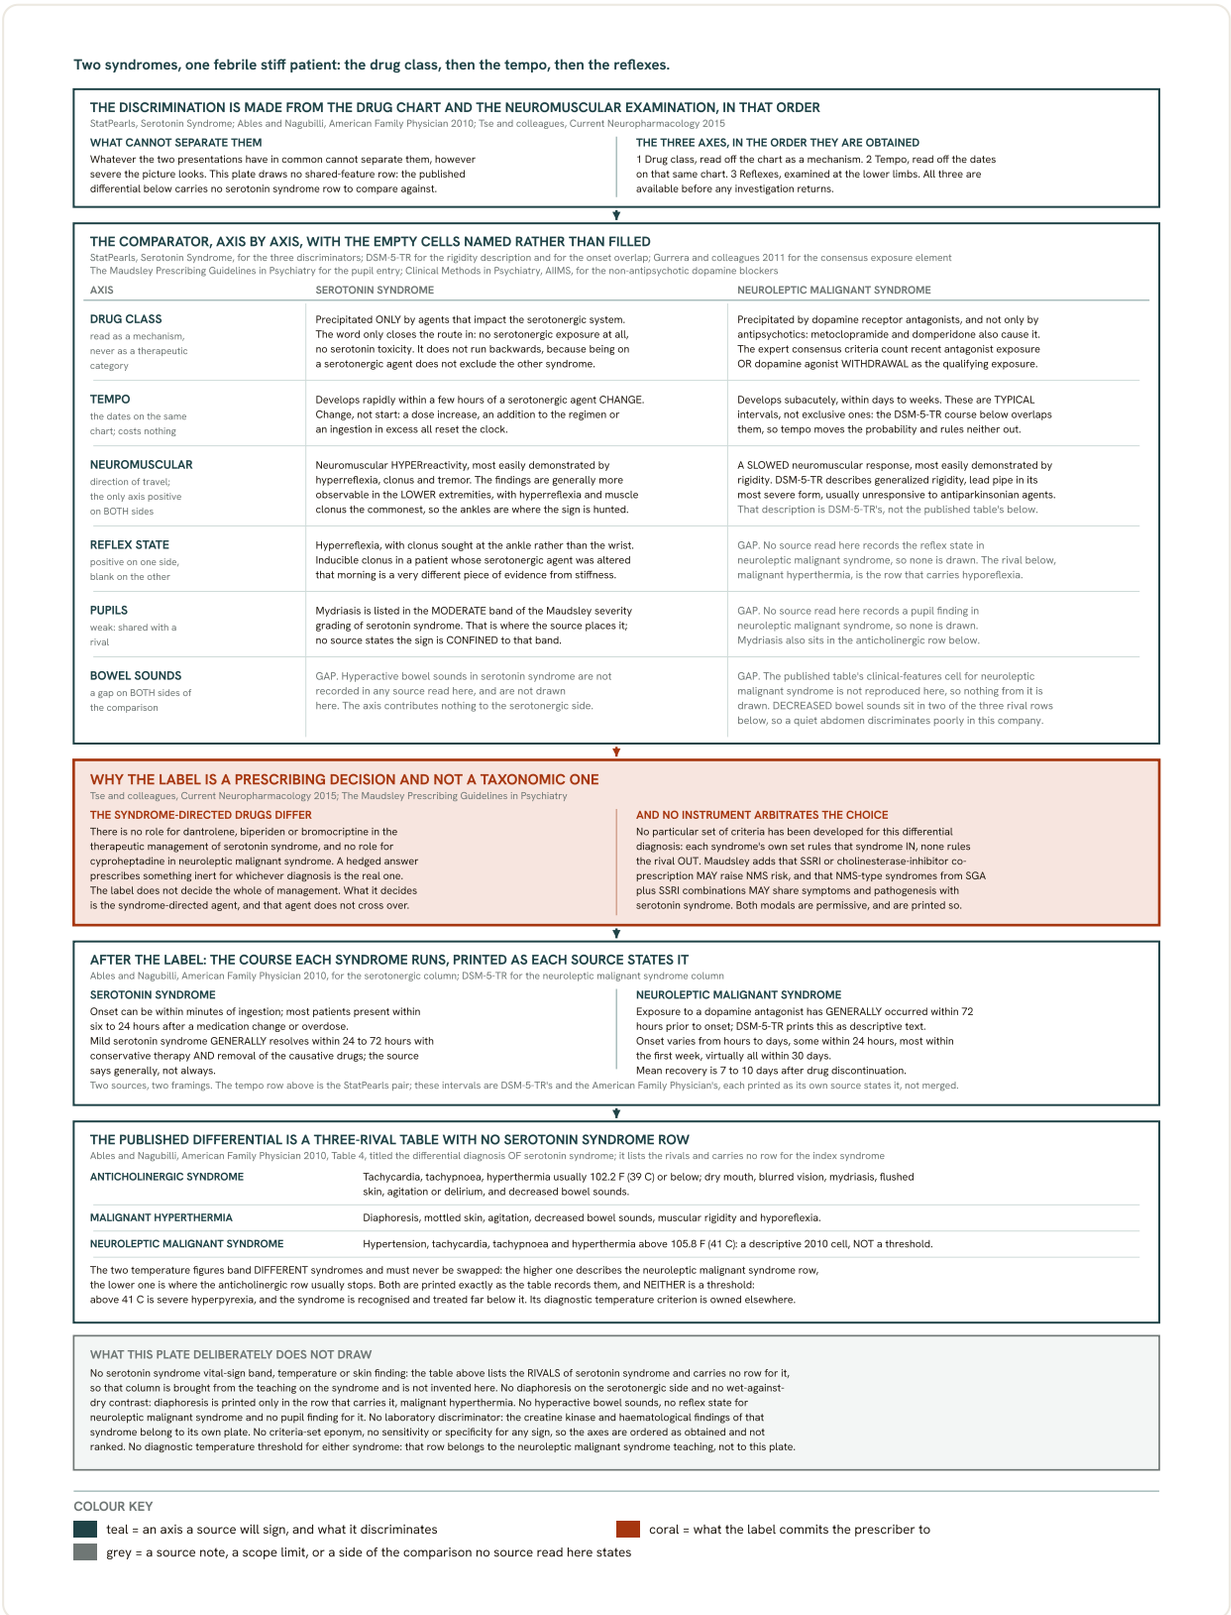


Fig 37.3 Serotonin syndrome against neuroleptic malignant syndrome: the axes that separate them, and the axes that are blank on one side. The three discriminators are obtained in order, drug class then tempo then reflexes, and only the neuromuscular axis returns a positive finding on both sides. Pupils, 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

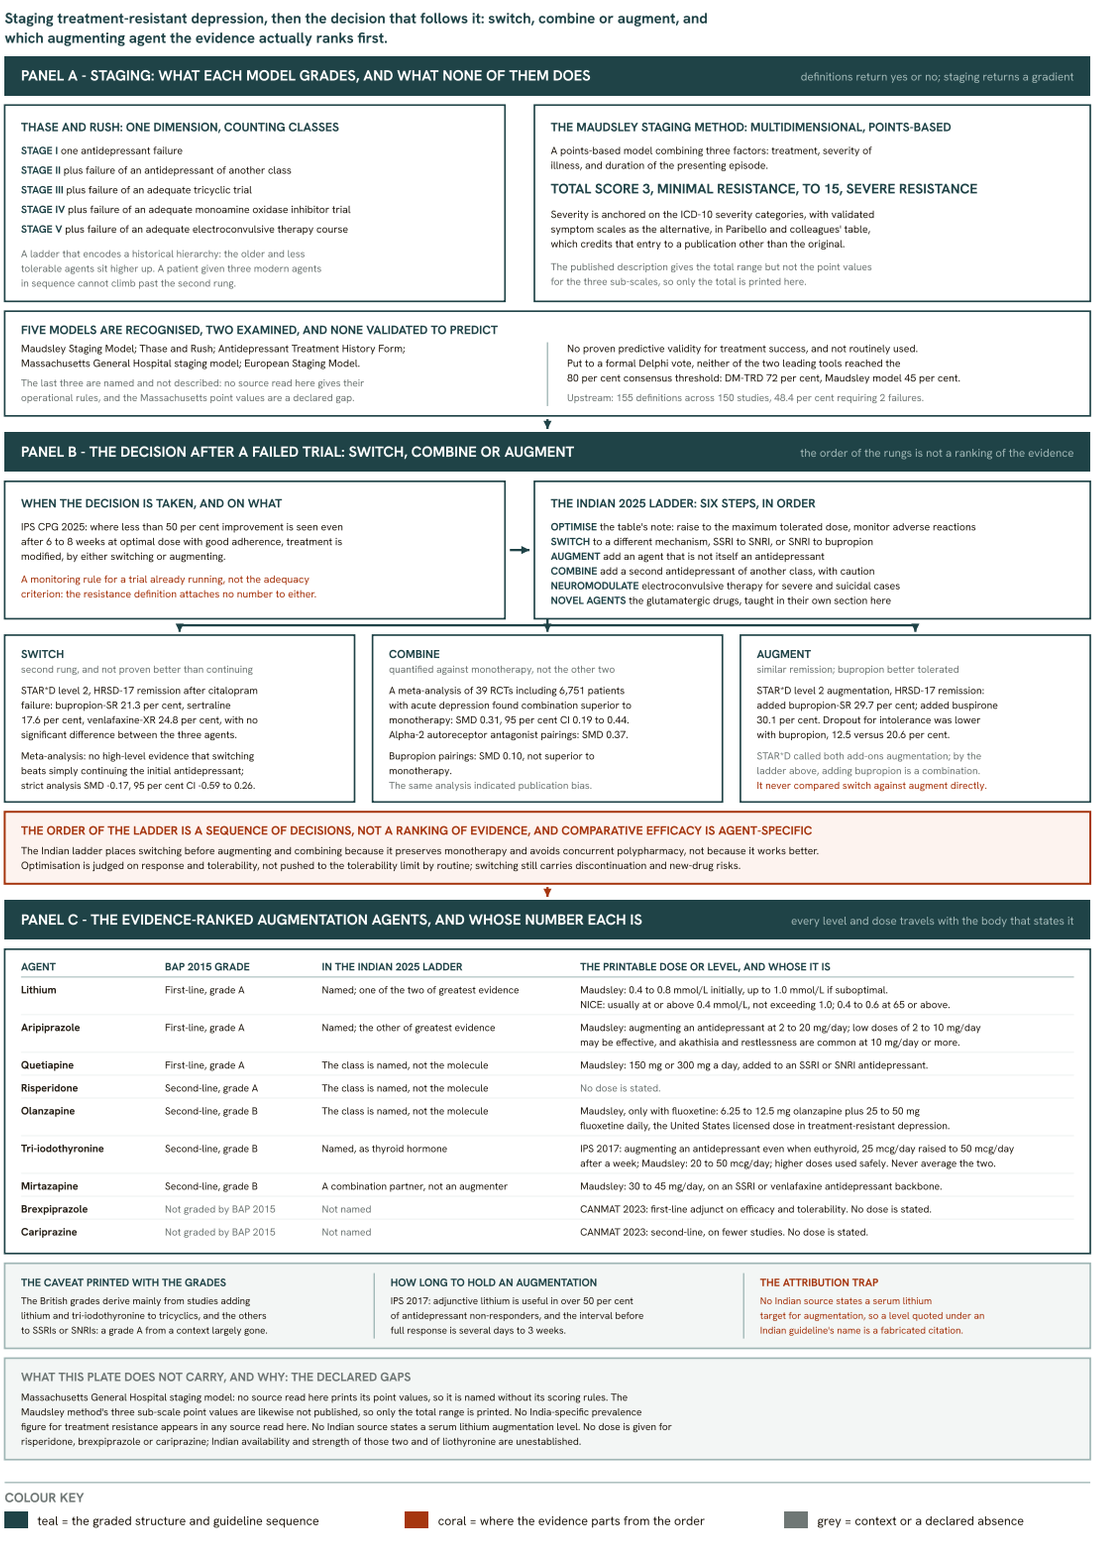


Fig 37.4 Staging treatment-resistant depression, then switching, combining or augmenting. Two of five recognised models are examined: Thase and Rush counts failed classes across five stages, the Maudsley method scores treatment, severity and duration to 3 to 15. Neither is validated to predict, and at 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

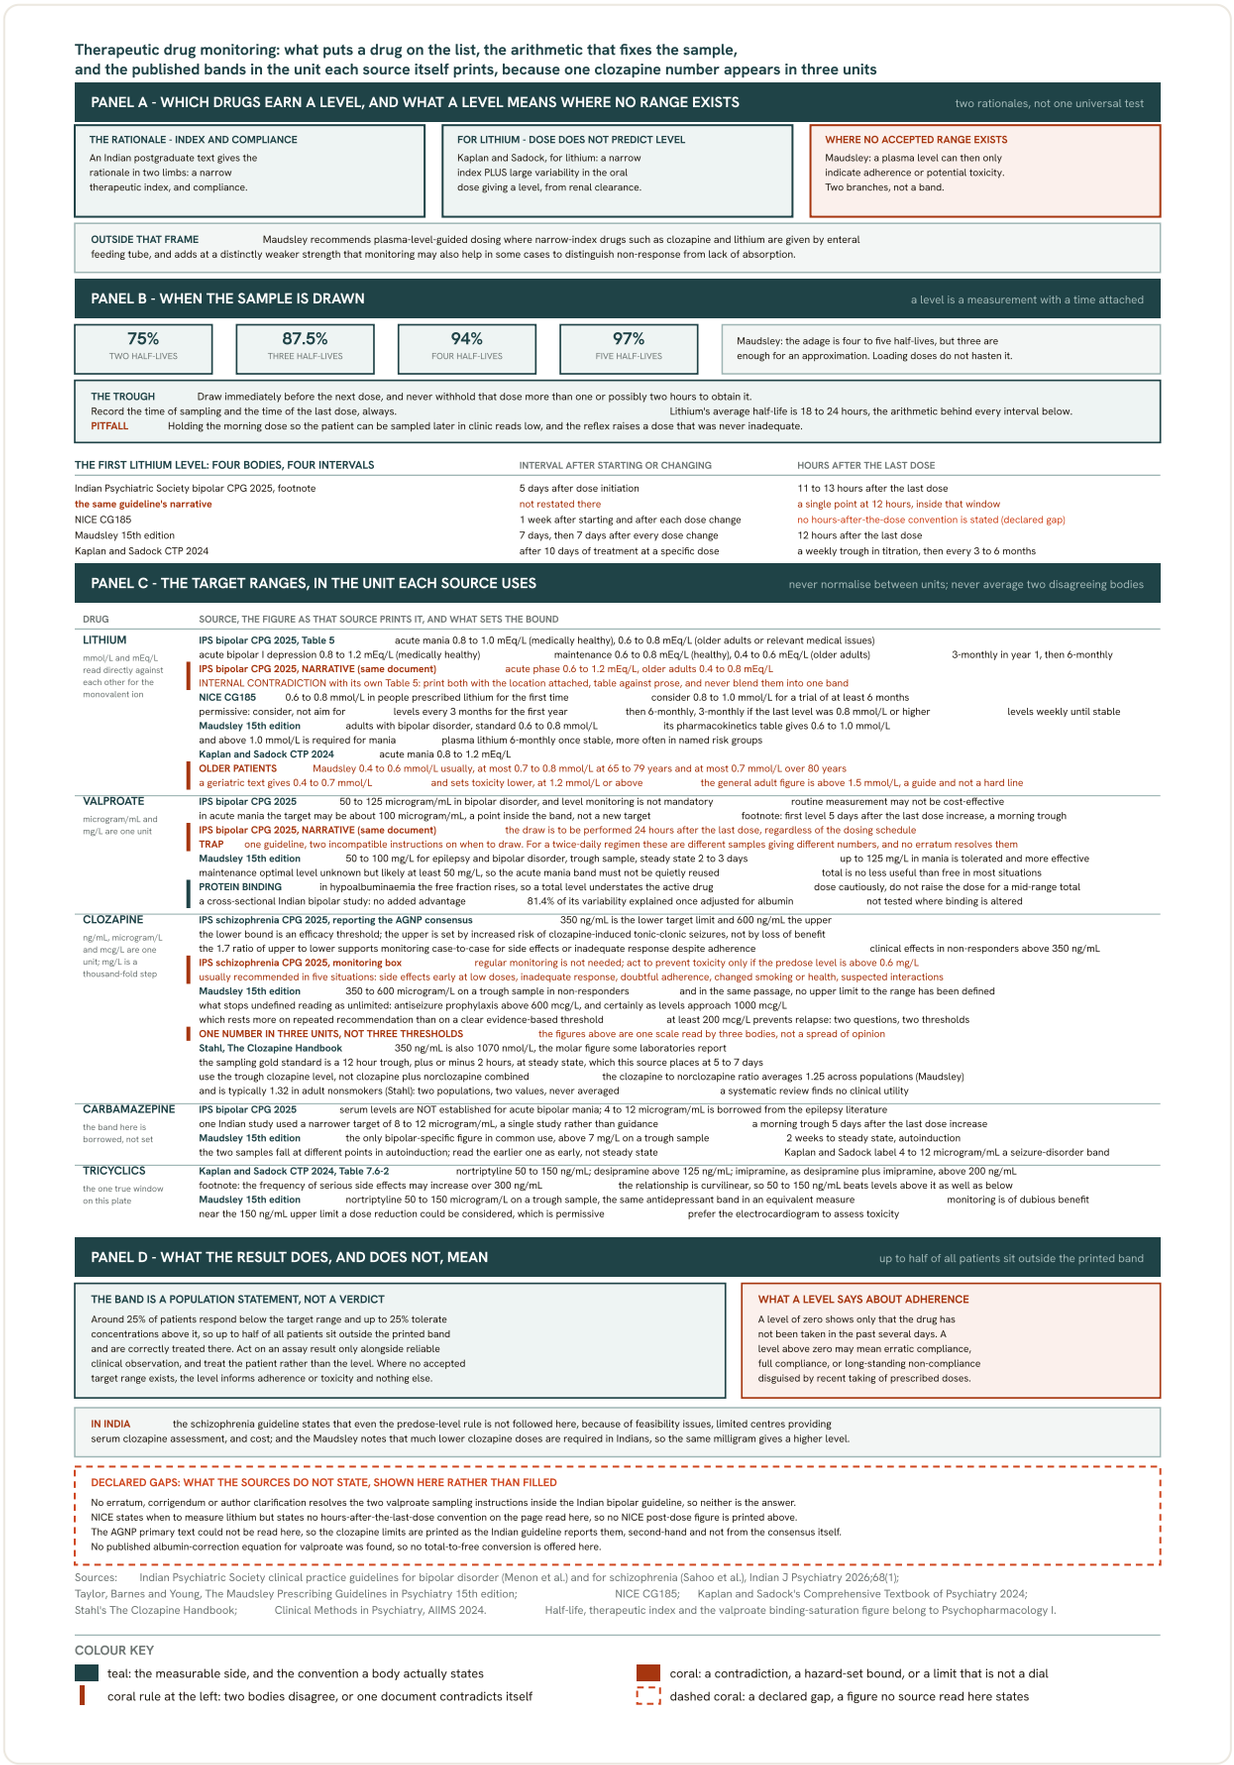


Fig 37.5 Therapeutic drug monitoring: which drugs earn a level, when the sample is drawn, and the bands as each source prints them. The lithium bands differ across four bodies, the Indian bipolar guideline contradicts

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

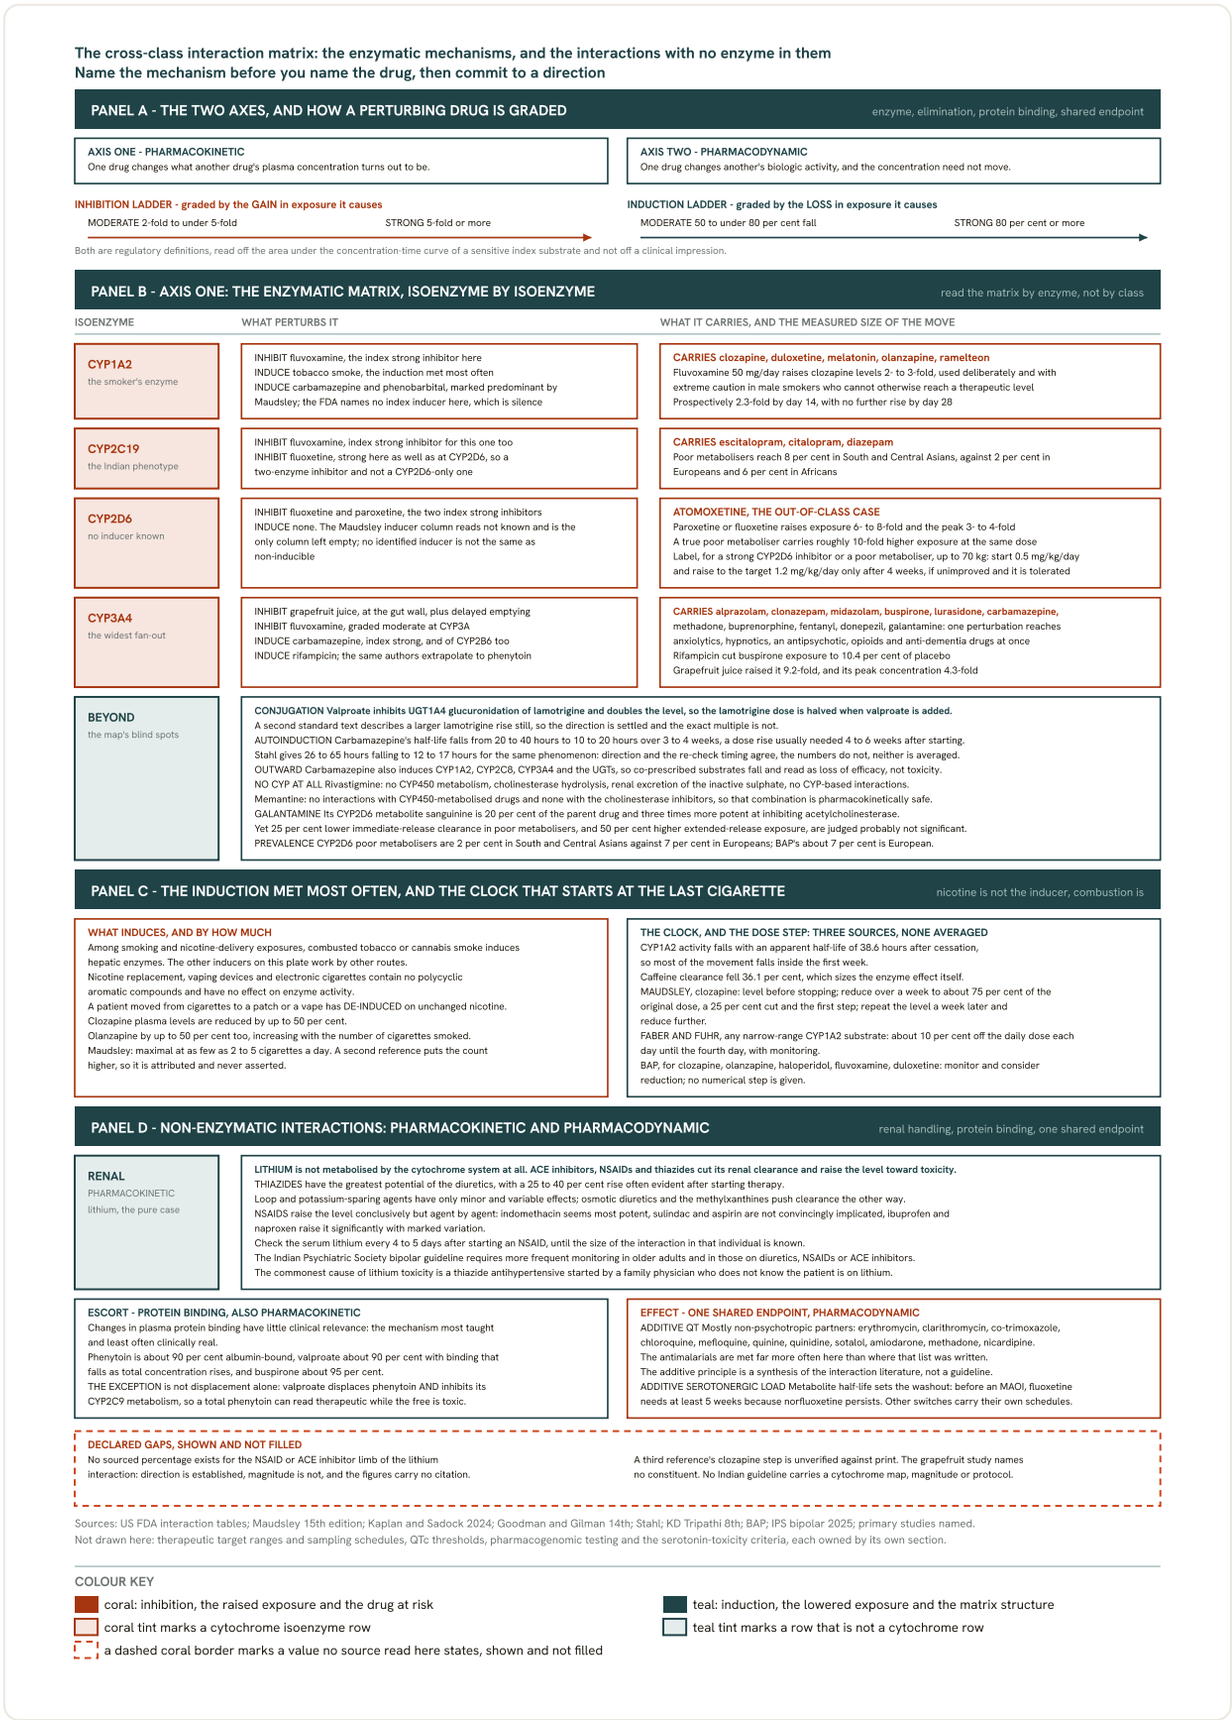


Fig 37.6 The cross-class interaction matrix, on both of its axes. The enzymatic half is read isoenzyme by isoenzyme: what inhibits it, what induces it, what it carries, and the measured size of the move, with conjugation, autoinduction, the agents that have no cytochrome involvement at all, and the population 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.

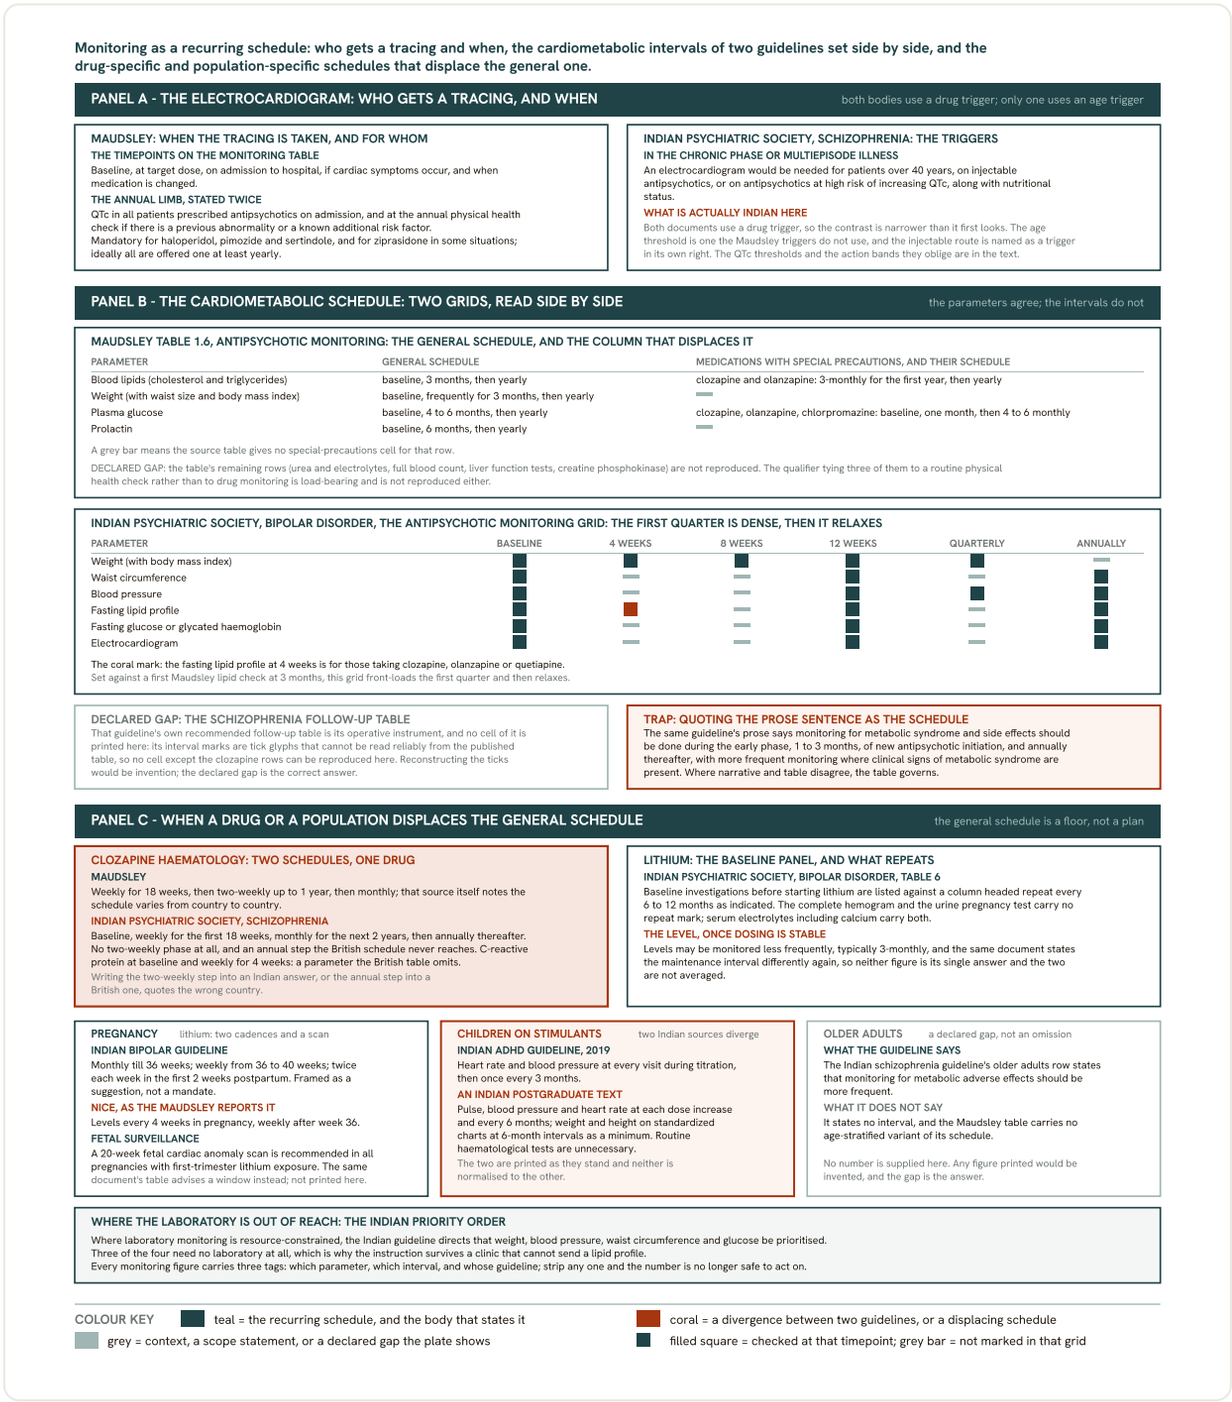


Fig 37.8 The recurring monitoring schedule: who gets a tracing and when, the cardiometabolic intervals of two guidelines set side by side, and the drug-specific and population-specific schedules that displace the general one. The Maudsley table is built around one early check and an annual rhythm, and 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.

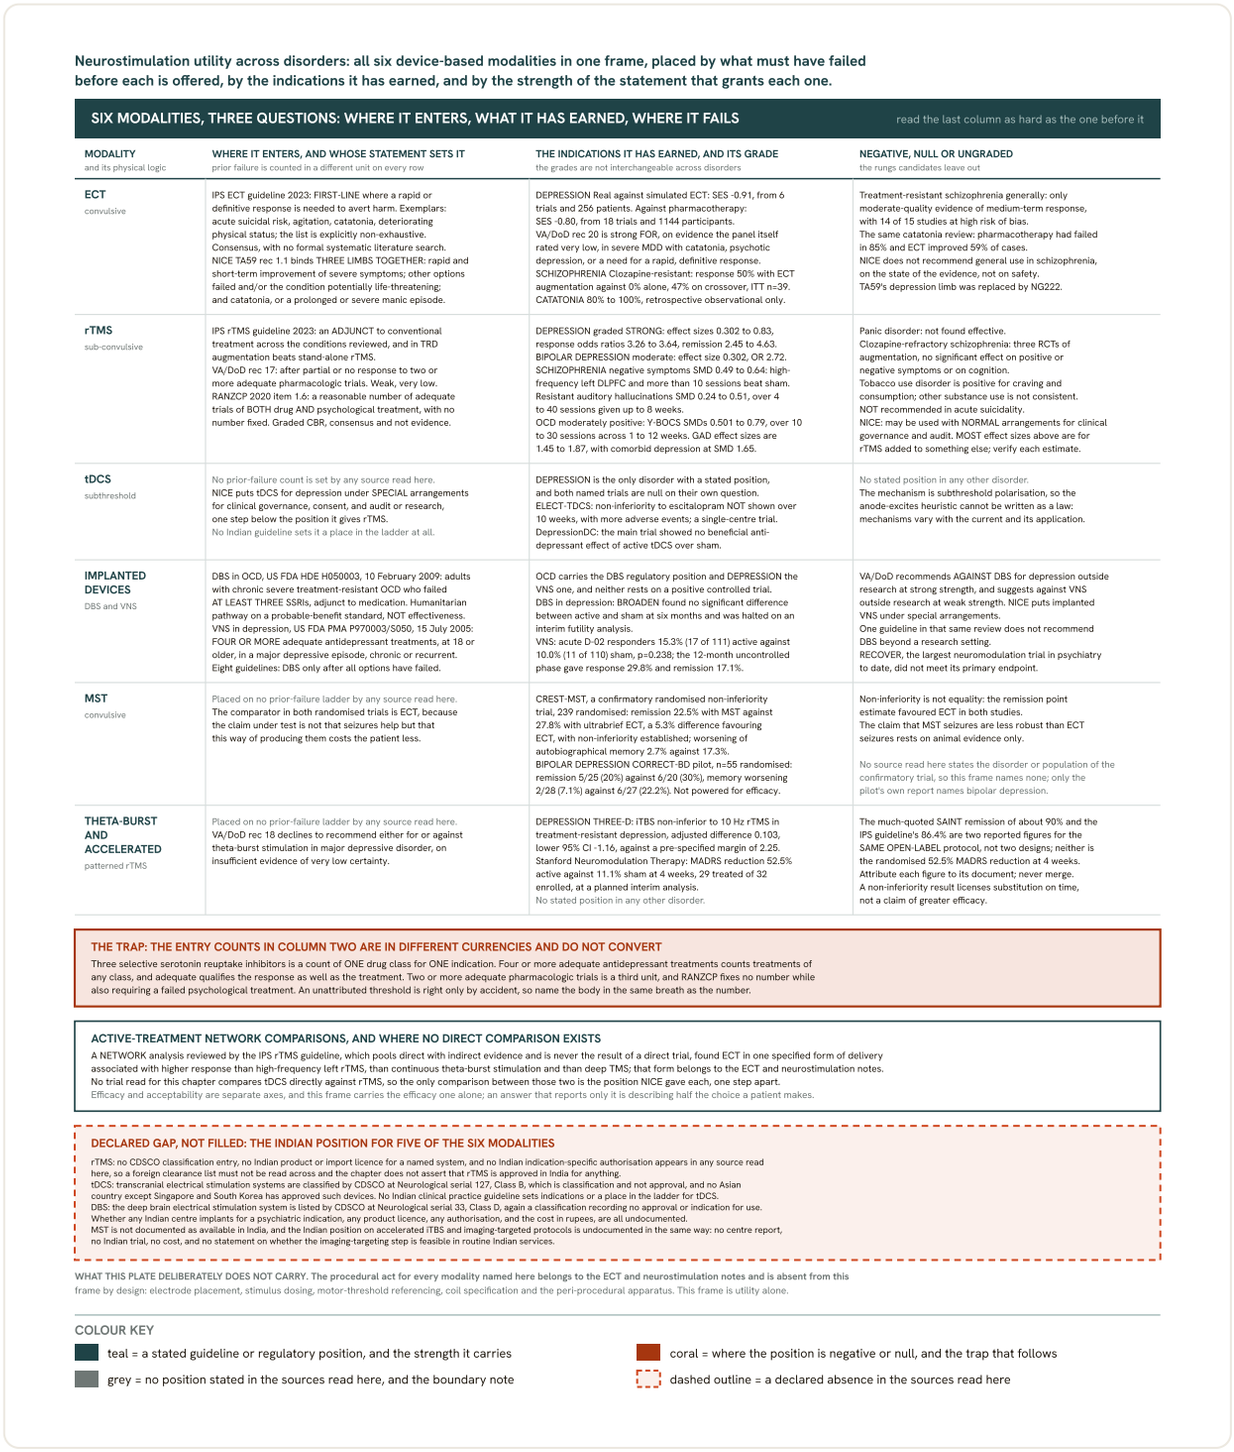


Fig 37.9 Neurostimulation utility across disorders: six modalities placed by entry point, by indication, and by the strength of the statement that grants each one. The entry column is a set of thresholds and not a sequence, because the counts are in different currencies and do not convert: no failure 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	4 or more	8	1
SSRIS FAILED, US OCD PATHWAY	ANTIDEPRESSANT TRIALS, US VNS INDICATION	GUIDELINES: ONLY AFTER ALL ELSE FAILS	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

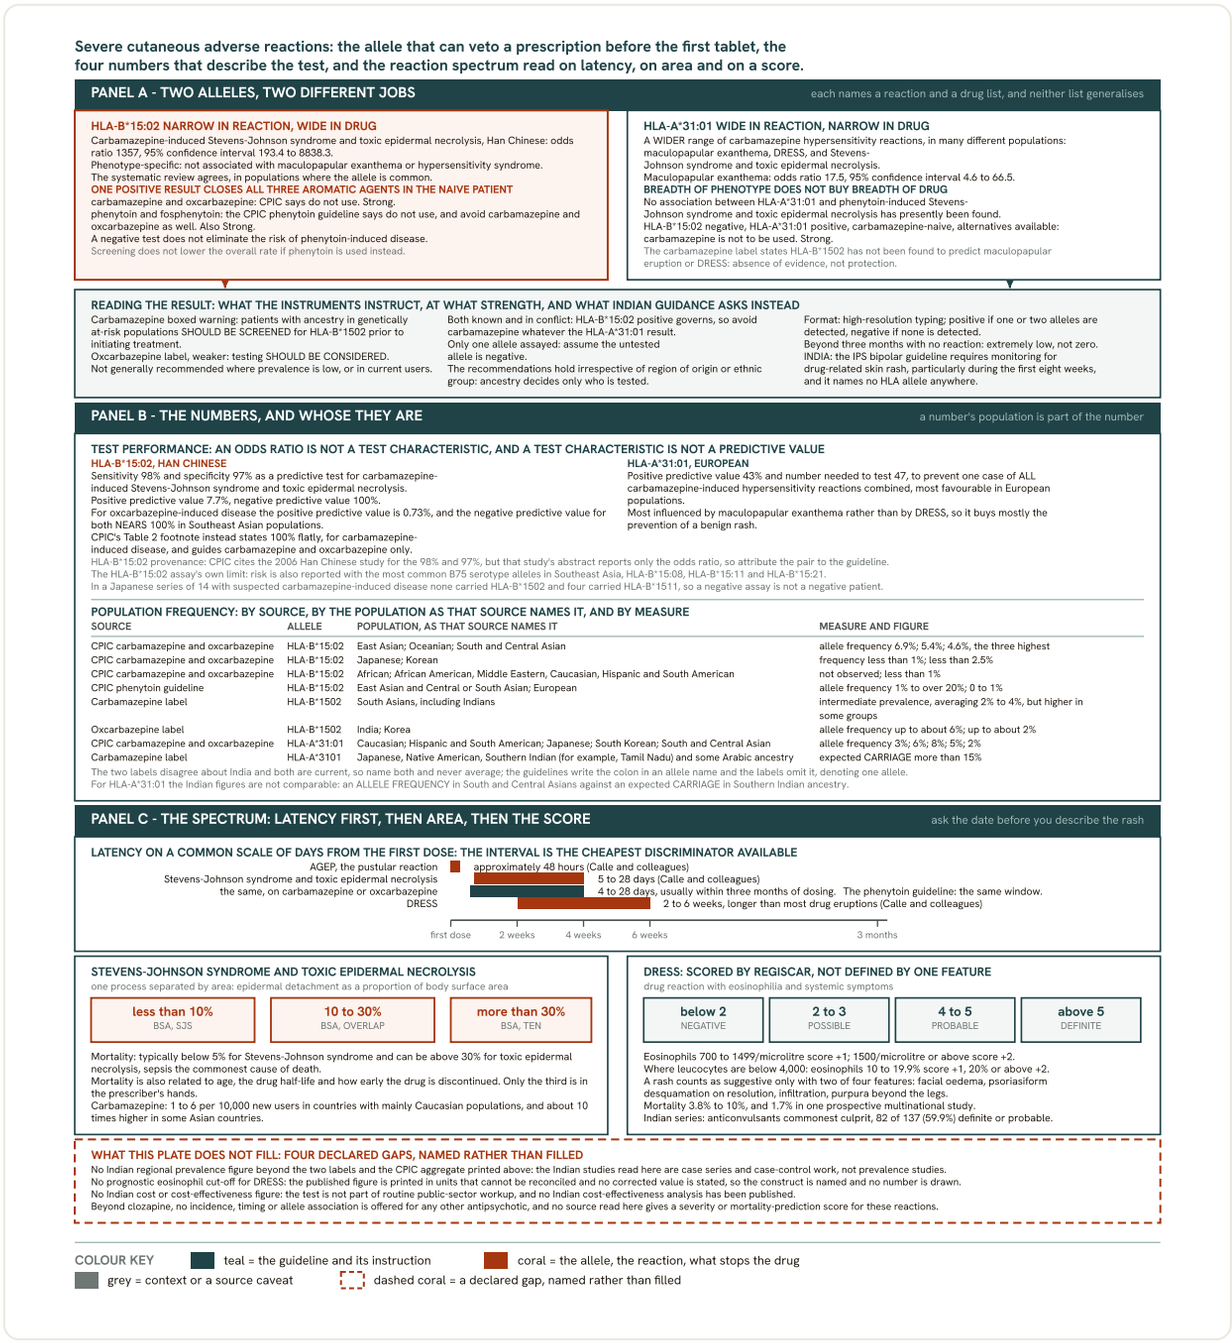


Fig 37.10 Severe cutaneous adverse reactions: allele to drug to population to test decision, then the spectrum the test exists to prevent. Two alleles do two different jobs. HLA-B*15:02 is narrow in reaction and wide in drug, carrying an odds ratio of 1357 for carbamazepine-induced Stevens-Johnson syndrome 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 non-adherence	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)

Index

Every entry jumps to its page. Numbers are the page of first mention.

#

5-HT2A antagonists 27

A

absolute neutrophil count 73
accelerated protocol 343
accumulation 33
acetylcholinesterase 176
activating antidepressant 110
active metabolite 70
activity score 284
Acute akathisia 62
Acute dystonia 49
acute on chronic 233
adequate antipsychotic trial 262
adequate dose and duration 238
adjunct 179
adverse drug event 275
adverse drug reaction 275
adverse-effect signature 12
agent that does not increase prolactin 88
AGEP 372
agonist spectrum 25
agranulocytosis 72
allele frequency 151
Allosteric modulation of nicotinic receptors 176
alpha-1 adrenergic 19
alpha-1 adrenergic antagonism 21
alpha-1 adrenergic receptors 19
alpha-2-delta protein 175
alternative pathway 36
amantadine 217
anterior pre-optic nucleus 216
anticholinergic burden 20
anticholinergic syndrome 225
anticholinergic toxidrome 231
anticipation 202
anticipatory monitoring 391
antimuscarinic blockade 20
antipsychotic polypharmacy 216
antipsychotic-induced cardiomyopathy 299
aromatic anticonvulsants 374
attribution 21
atypical 15
atypical antipsychotic 218
augmentation 204
autobiographical memory 334
autoinduction 15
azapirone 173

B

benign ethnic neutropenia 73
benzodiazepine receptor agonists 166
benzodiazepine-sensitive 160
bidirectional 299
binding profile 12
Binding promiscuity 188
Bioavailability 31
biogeographic group 284
biperiden 217
Bocquet criteria 378
boundary handover 397
boxed warning 155
bromocriptine 217
butyrophenones 48
butyrylcholinesterase 176

C

catatonia 217
CET score 383
cholinesterase inhibitors 175
Chronic akathisia 64
chronic lithium toxicity 135
chronic low white cell count 128
circuit therapy 350
Class I antiarrhythmics 116
class-defining hazard 195
Clearance 16
clonus 179
clozapine-refractory schizophrenia 334
clozapine-resistant schizophrenia 332
coarse tremor 128
cold retrieval 196
complete lactation failure 318
concentration to dose ratio 289
conjugation to inactive glucuronide metabolites 153
constitutive activity 26
covert non-adherence 263
creatine kinase 217

creatinine clearance 299
creatinine creep 304
cyclopyrrolone 161
CYP2D6 37
CYP2D6*10 allele 118
CYP3A subfamily rather than for CYP3A4 alone 37
cyproheptadine 212

D

dantrolene 217
de-esterification 182
de-induction 40
dechallenge 277
deep brain stimulation 334
deep TMS 334
Delayed onset of action 99
Demethylation 114
dextromethorphan-bupropion 254
diaphoresis 209
difficult-to-treat depression 245
DiHS 379
diphenylbutylpiperidines 48
Discontinuation symptoms 100
disproportionality analysis 281
dissociative symptoms 259
dopamine antagonist 173
dopamine serotonin antagonist 52
dopamine transporter inhibition 102
dose-dependent dual action 108
DRESS 113
dual noradrenaline and dopamine reuptake inhibitor 96
dysgeusia 161

E

effect change 27
effect duration 181
Efficacy 34
ELECT-TDCS 334
electroconvulsive therapy 222
elemental lithium 130
elimination half-life 13
emergency act 203
Enzyme inhibition 15
enzyme-inducing anticonvulsant 154
ester prodrug formulated as an aqueous nanosuspension 78
euthyroid 243
exposure change 391
extrapyramidal nervous system 54

F

familiarity heuristic 195
fast-off property 56
first generation 96
First-order kinetics 31
first-pass metabolism 39
flumazenil 162
free fraction 34
full agonist 25
full gradual escalation regimen 158
fulminant hepatic failure 128
futility analysis 334

G

gastrointestinal hypomotility 76
generalised anxiety disorder 341
generalised bullous fixed drug eruption 382
generalised erythematous rash 150
glucuronidation 128
Glycogen synthase kinase 3 129
graded dose-response 43
growth suppression 320

H

handover boundary 16
Hardy-Weinberg equilibrium 285
hepatic biotransformation 385
hepatic encephalopathy 299
hepatic metabolising capacity 298
HIFU 360
high-tyramine meal 120
histaminergic H1 blockade 20
histone deacetylase 128
HLA-A*31:01 286
HLA-B*15:02 286
human leukocyte antigen (HLA) variant B*15:02 150
Humanitarian Device Exemption 351
Hunter Serotonin Toxicity Criteria 211
hyperbolic 161
hyperreflexia 209
hypertonicity 209
hypoalbuminaemia 270
hypokalaemia 323
hyponatraemia 102

I

iatrogenic differential 396
iatrogenic phenocopy 388
idiosyncratic 24
imidazopyridine 161

impulse control disorders 299
inadequate response 238
index substrate 292
indirect sympathomimetic 122
induction phase 257
Information Component 281
inhibition of dopamine uptake 96
inhibits neuronal sodium channels 152
Inositol monophosphatase 129
intention or postural tremor 142
inter-ictal psychosis 309
interface of the gamma-2 and alpha subunits 159
interleaved retrieval 194
intermediate metaboliser 38
intrinsic activity 190
inverse agonist 25
ionotropic receptor 22
irreversible inhibitors 122

L

label-driven prescribing 396
laryngospasm 259
latency 63
law of mass action 56
lead-pipe 216
leucopenia 151
licarbazepine 146
linezolid 208
loading doses 268
long QT syndrome 323

M

magnetic seizure therapy 338
major adverse cardiac events 336
major congenital malformations 313
malignant catatonia 228
malignant hyperthermia 225
Maudsley Staging Method 243
medical device 349
melatonin receptor agonist 111
melatonin receptor agonists 161
membrane potential 345
membrane stabilisation 116
mesocortical dopamine pathway 55
mesolimbic dopamine pathway 54
metabolic ratio 71
metabolic syndrome 326
metaboliser phenotype 38
microspheres 77
moclobemide 120
mode of action 12
monitoring proxy 395
monohydroxy derivative 149
monotherapy 128
multimodal 110
muscarinic cholinergic affinity 102
muscarinic cholinergic receptors 19
Myocarditis 69
myoclonus 179

N

narrow therapeutic index 46
negative predictive value 288
neonatal adaptation difficulties 316
neonatal bleeding diathesis 314
nephrotoxic 299
network meta-analysis 86
neural tube defects 128
Neuroleptic-induced dysphoria 88
neuromuscular hyperreactivity 225
Neuroscience-based Nomenclature 52
new-onset mania 348
nigrostriatal dopamine pathway 54
NMDA receptor antagonism 19
non-inferior 334
non-inferiority 334
non-invasive brain stimulation 344
non-response 38
nootropic 179
noradrenergic and specific serotonergic antidepressant 109
norclozapine 71
Norfluoxetine 102
number needed to test 370

O

occupancy threshold 56
off-target action 13
offset 13
oil vehicle 78
orexin receptor antagonists 161
orthostatic hypotension 90

P

partial 5-HT1A agonist 173
partial agonist 25
partial response 29
patch testing 383
pentamer 159
persistent pulmonary hypertension of the newborn 316
pharmacodynamic 42
pharmacodynamic interaction 290

Pharmacodynamics 42
pharmacokinetic 11
pharmacokinetic interaction 36
pharmacological discriminator 187
pharmacovigilance 267
phasic signalling 178
phenoconversion 189
phenothiazine 48
physostigmine 232
piperazine phenothiazines 48
piperidine phenothiazines 48
polycyclic aromatic compounds 40
polycyclic aromatic hydrocarbons 40
poor metaboliser 36
poor metabolisers 36
positive allosteric modulator 19
positive predictive value 370
post-ictal psychosis 299
Post-injection delirium sedation syndrome 79
Potency 26
pregnancy prevention programme 128
premarket approval 351
presystemic extraction 35
Priapism 19
procedural act 203
proconvulsive 299
prodrug 78
prolactin 19
prolactin neutral 88
protracted withdrawal syndrome 159
Pseudo-irreversible inhibitors 176
Pseudoparkinsonism 62
pseudoresistance 262
psycho-pharmacogenomic studies 289
psychosurgery 355
pyrazolopyrimidine 161

Q

QTc 292
quantal dose-response 43

R

racemic ketamine 254
randomised withdrawal 258
receptor binding profile 12
Receptor cleanliness 100
rechallenge 73
reference schedule 156
RegiSCAR 372
release mechanism 182
renal elimination 290
renal prostaglandins 137
retrieval card 193
retrograde amnesia 335
reverse inference 187
reverse retrieval 395
reversible inhibitor of MAO-A 123
risk allele 288
ritalinic acid 181
ryanodine receptor 217

S

sanguinine 177
Schedule III controlled substance 259
second generation 47
second-generation antipsychotics 314
secondary amines 114
selective noradrenaline reuptake inhibitor 184
separating axis 193
serious 44
serotonin toxicity 100
serotonin transporter 95
set point 28
severe neutropenia 69
sexual dysfunction 56
sham 334
signal 12
SILENT 25
silent antagonist 25
simple branched-chain fatty acid 140
simulated treatment 332
single-centre trial 334
slow-dose stimulants 183
slow-release 110
sodium and hydrogen exchanger 3, or NHE3 130
sodium, potassium and chloride cotransporter NKCC2 136
spaced retrieval 197
stabilisers 15
staccato speech 231
staging models 238
Stanford Neuromodulation Therapy 334
steady state 13
Sternbach criteria 211
Stevens-Johnson syndrome 155
strongly antimuscarinic 20
sub-convulsive 334
subanaesthetic 255
subgenual anterior cingulate cortex 364
substituted benzamides 48
subthreshold intensity 366
sudden cardiac death 299
supersensitivity 64

sustained response 340
syndrome of irreversible lithium-effectuated neurotoxicity 138
systematic pharmacology 11

T

tACS 359
Tardive akathisia 64
tardive dyskinesia 19
tempo 39
tertiary amines 114
Thase and Rush model 244
Therapeutic drug monitoring 14
therapeutic drug monitoring 14
therapeutic index 14
therapeutic threshold 46
therapeutic window 47
theta rhythm 362
theta-burst stimulation 334
thioxanthenes 48
THREE-D 334
torsade de pointes 323
toxic epidermal necrolysis 288
transcranial magnetic stimulation 337
transporter preference 108
treatment-resistant depression 237
treatment-resistant schizophrenia 261
tri-iodothyronine 243
tRNS 359
trough 30
trough sample 30
tuberoinfundibular dopamine neurons 54
twelve hour trough 132
Type B 276
Tyramine 120

U

unbound fraction 141
uncompetitive open-channel NMDA receptor antagonist 178
Uncoupling 161

V

vagus nerve stimulation 334
ventral capsule/ventral striatum 351
visual hallucinations 299
voltage-dependent sodium channels 146
voltage-gated sodium channels 140
voltage-sensitive sodium channels 115
volume depletion 131
volume of distribution 30
Volume of distribution 30

W

waist circumference 326
washout 33
wide therapeutic index 45
within-class selection 192

Z

Zero-order kinetics 31

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

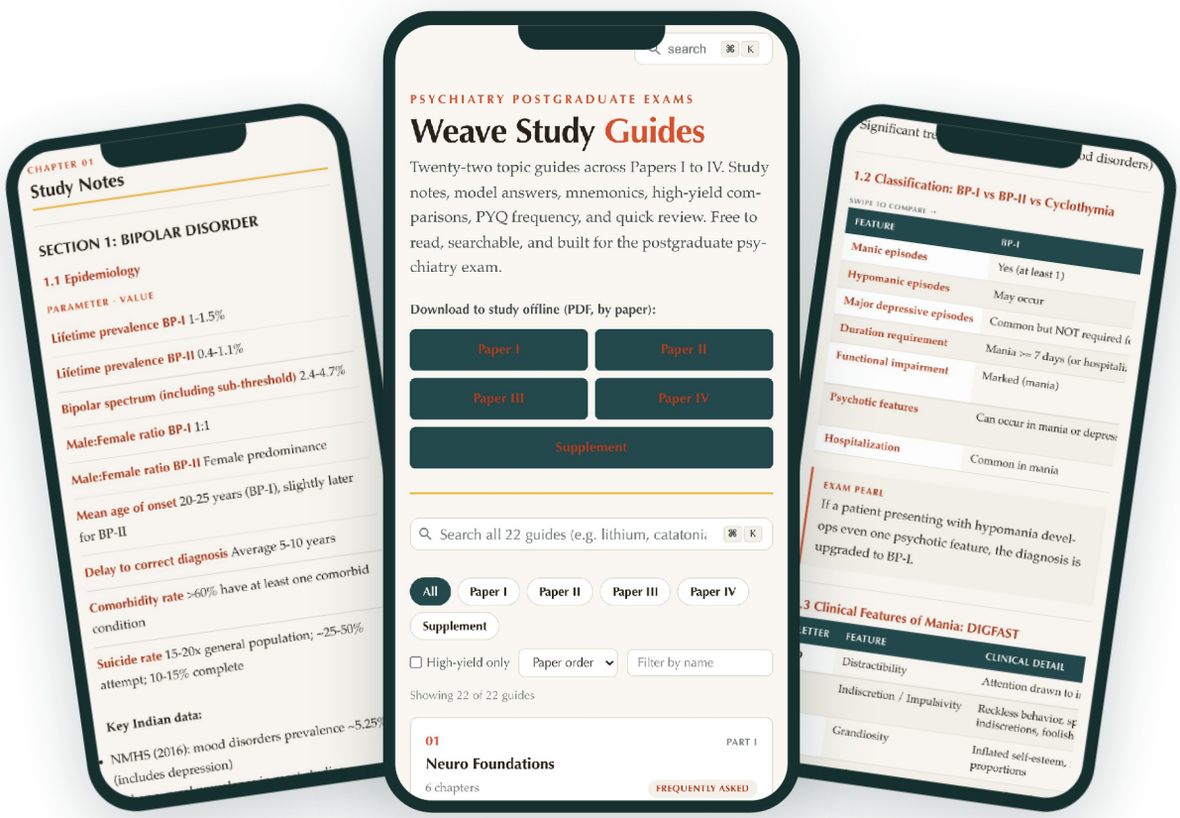
weave.clinic/library

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

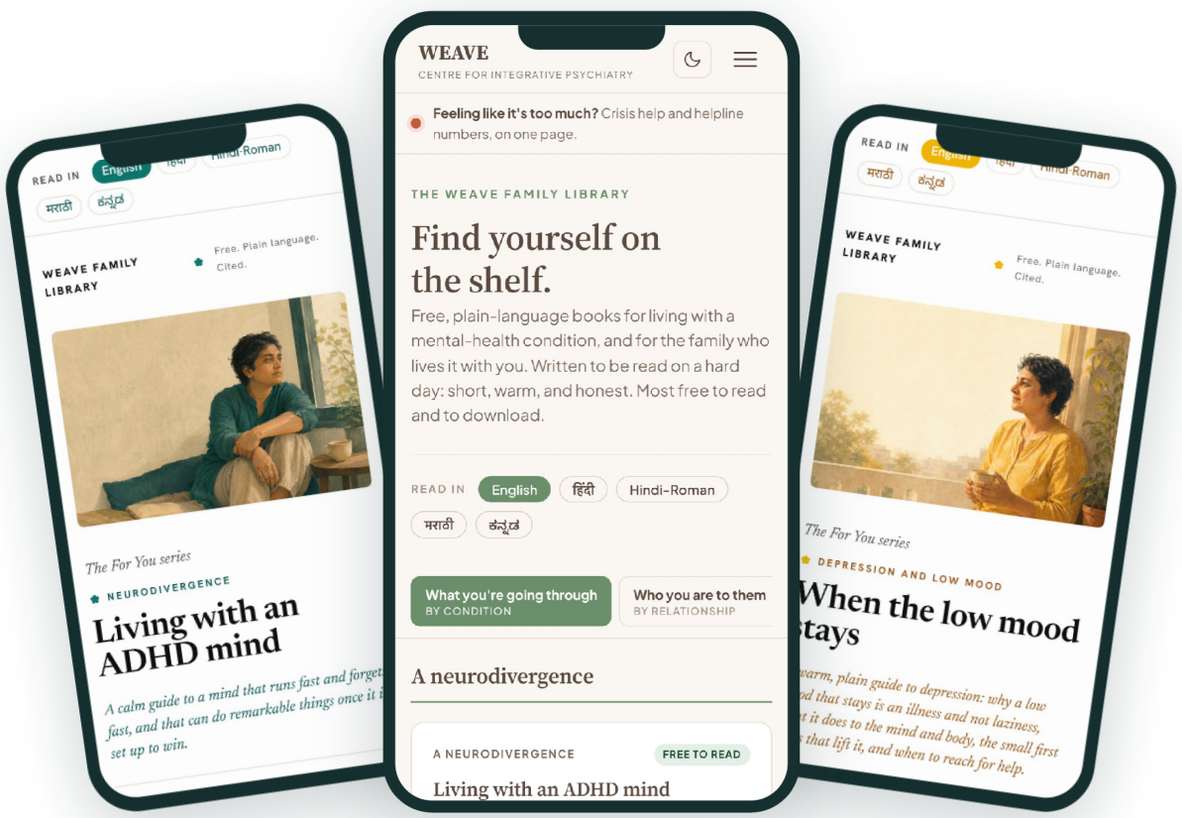
weave.clinic/study

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 · VOLUME FIVE

Volume Five of the fourth paper is psychopharmacology in full, across two chapters: the classes, the receptors, the adverse effects and the reasoning behind a first choice; then prescribing when it goes wrong or will not work, the emergencies, treatment resistance, the monitoring schedules, and the interactions that produce both. Two 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