WEAVE SPRINT . PAPER IV

Psychopharmacology I: principles and core drug classes

The psychotropic classes taught as pharmacology rather than as treatment: how a receptor profile predicts the adverse effects a patient will report, what half-life and steady state settle about when a plasma level means anything, the cytochrome map that turns a second prescription into an interaction, the antipsychotics and antidepressants compared inside their own classes rather than against placebo, lithium as the archetype of a narrow therapeutic index, the anticonvulsant mood stabilisers and the constraints that govern them, and the benzodiazepines, hypnotics and stimulants read off their receptor subtypes.

CONTENTS

Dr. Wilfred D'souza . Dr. Niharika Reddy

WEAVE . CENTRE FOR INTEGRATIVE PSYCHIATRY

When therapy is part of the plan, there is somewhere to refer.

For clinicians who assess, prescribe or manage care but do not provide psychotherapy themselves. Trijog gives patients a route to online therapy and in-person sessions in Mumbai.

A CLINICIAN REFERRAL ROUTE

Refer to Trijog for therapy.

Your patient can contact Trijog directly. Our team can help them choose online or in-person care and identify a therapist based on their needs and preferences. The patient decides whether to book and whether the fit is right.

booking.trijog.com · 1800-2027-180 · the Trijog app

Online therapy	In person in Mumbai	Across ages	Across relationships
For patients who prefer to meet remotely.	Powai, Bandra and Prabhadevi.	Adults, children and teenagers.	Individuals, couples and families.

What the referral route keeps clear

01
Your role

You remain the referring clinician. A therapist takes on the psychotherapy work.

02
The fit

Concern, age, language, format and preferred way of working can all shape the match.

03
The setting

Your patient can meet online, or in person at one of Trijog's Mumbai clinics.

04
The coordination

If your patient wants it, their clinicians can share what is needed to keep care aligned.

Weave, powered by Trijog

PSYCHIATRY VERTICAL

Weave is the psychiatry vertical of Trijog. Weave Sprint is part of its Psychiatry Outreach Programme, supporting psychiatry residents and clinicians with practical learning and a clear route into Trijog's wider care network.

The notes in your hands are part of that programme. Therapy through Trijog can sit alongside psychiatric care when the patient and clinicians choose.

You do not have to provide therapy yourself to make it part of good care.

Share this route with a patient

Online therapy or in-person care in Mumbai.
booking.trijog.com · 1800-2027-180

Keep consent visible

Any coordination between clinicians happens with the patient's permission. The patient remains free to choose whether to book.

Routine referrals only. This is not an emergency route. For urgent risk, follow your local emergency and safeguarding pathway.

About this insert. Weave Sprint is educational outreach from Weave, the psychiatry vertical of Trijog. Therapy is provided through Trijog.

Contents

Concept / rule Key number Trap

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.

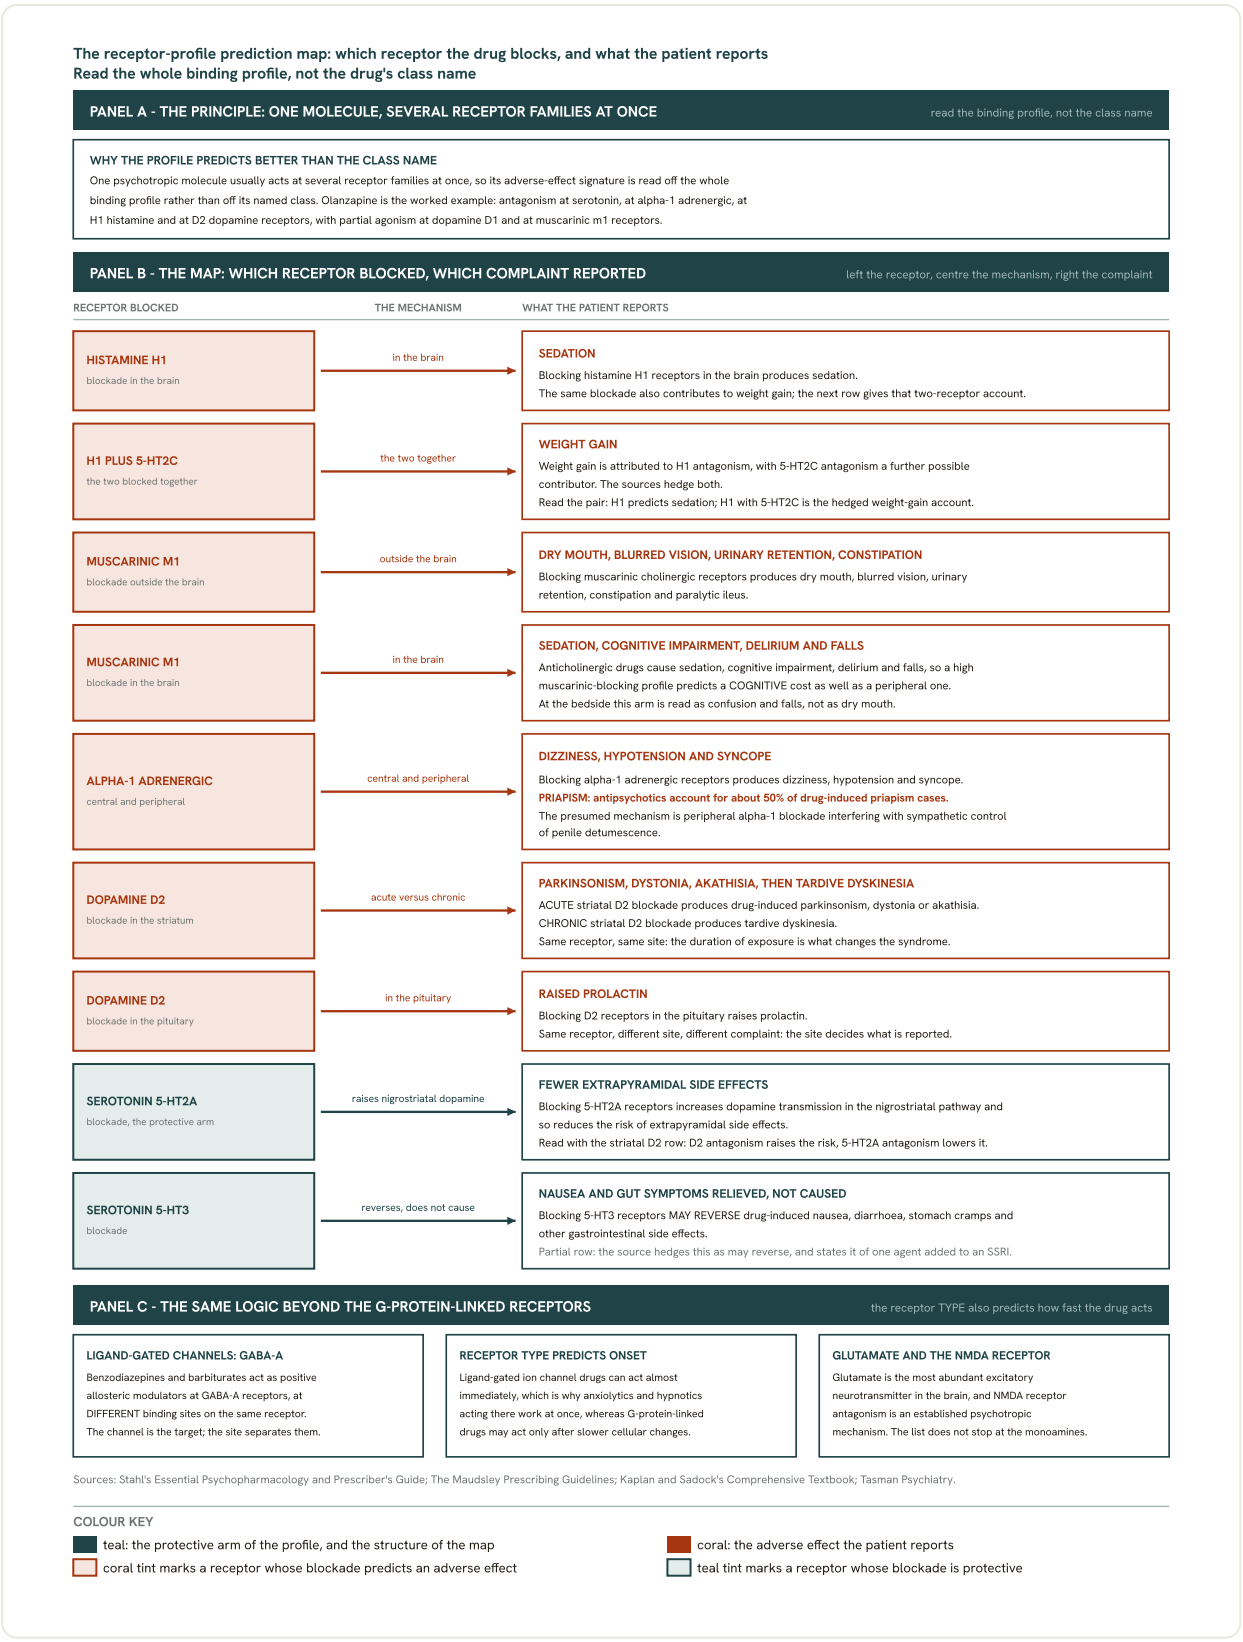

Fig 36.1 **The receptor-profile prediction map.** The chapter's organising figure: the adverse effect a patient reports is read off the drug's whole binding profile, not off its class name. H1 blockade predicts sedation, H1 with 5-HT2C weight gain, muscarinic blockade a peripheral and a cognitive cost, 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.

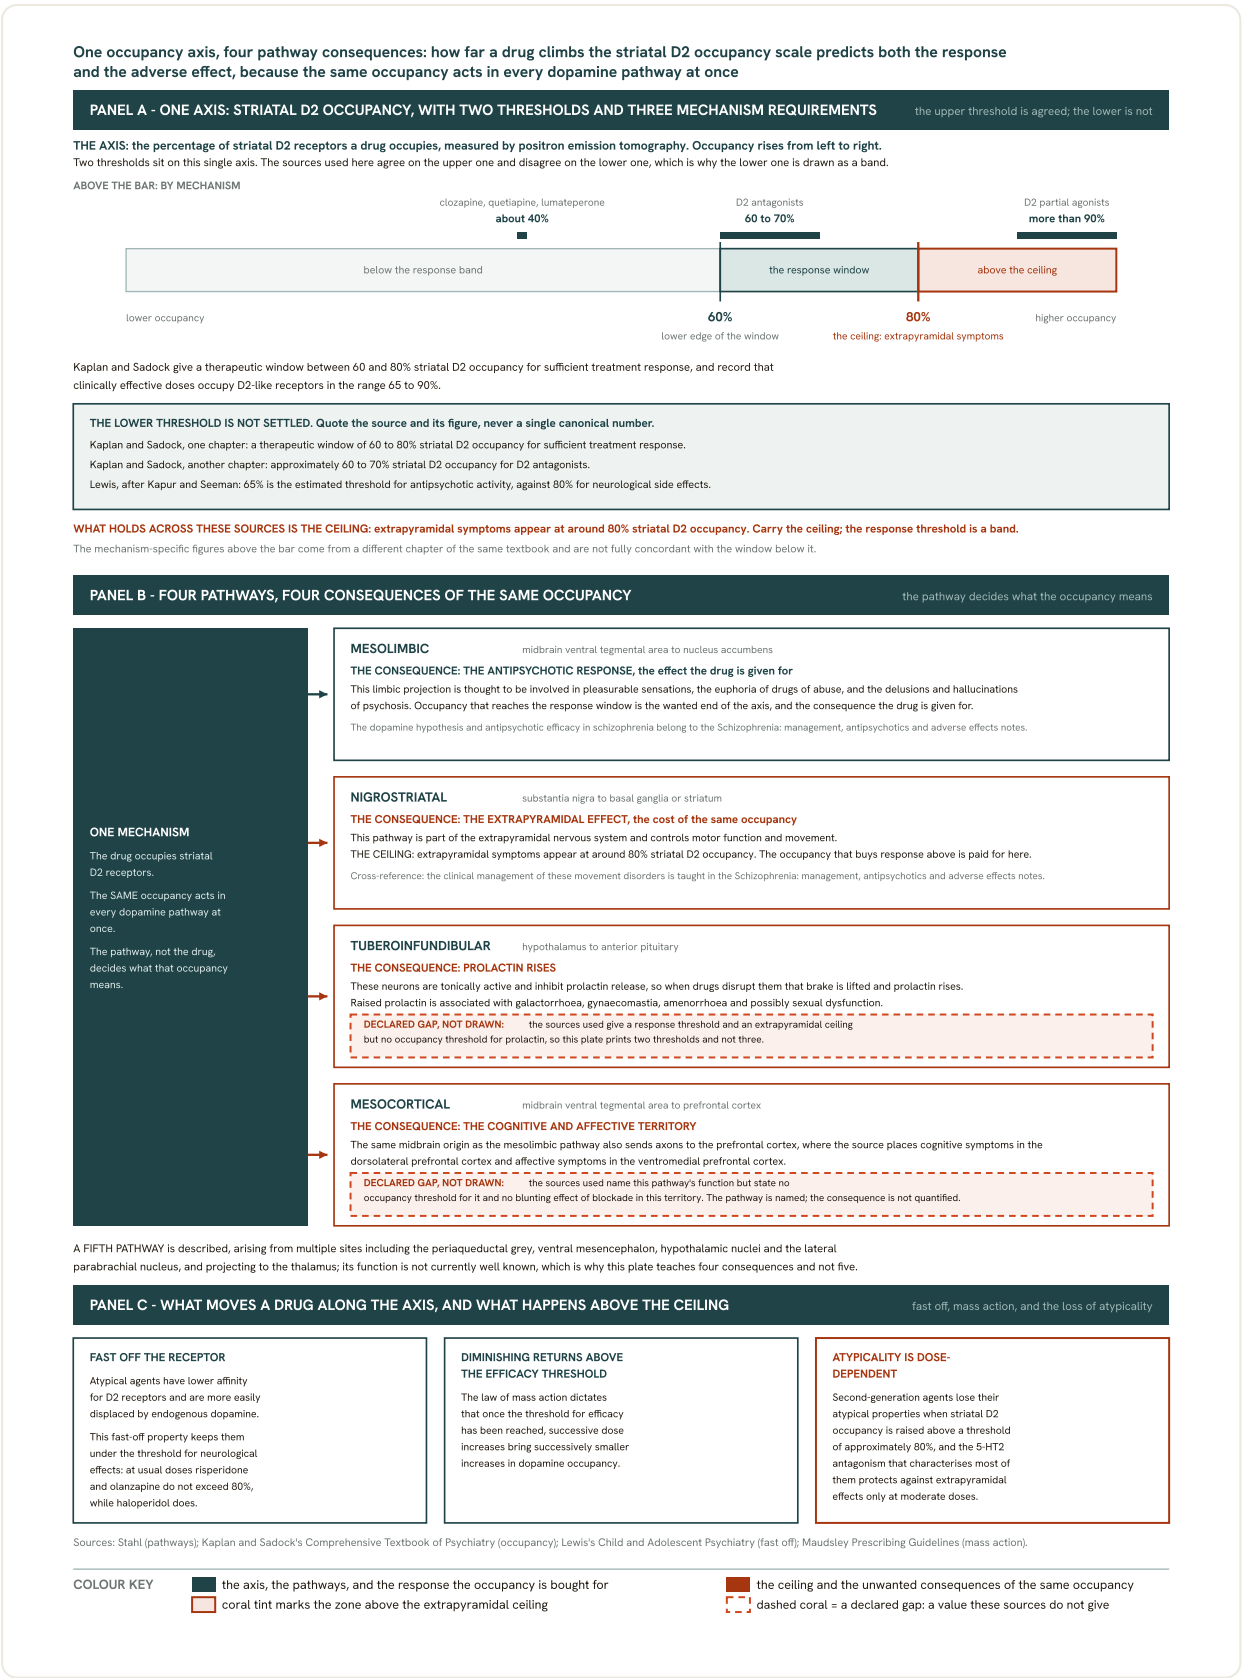

Fig 36.2 One occupancy axis, four pathway consequences. How far a drug climbs the striatal D2 occupancy scale predicts both its response and its adverse effects: the sources here agree on a ceiling of around 80% above which extrapyramidal symptoms appear, and disagree on the response threshold, placing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

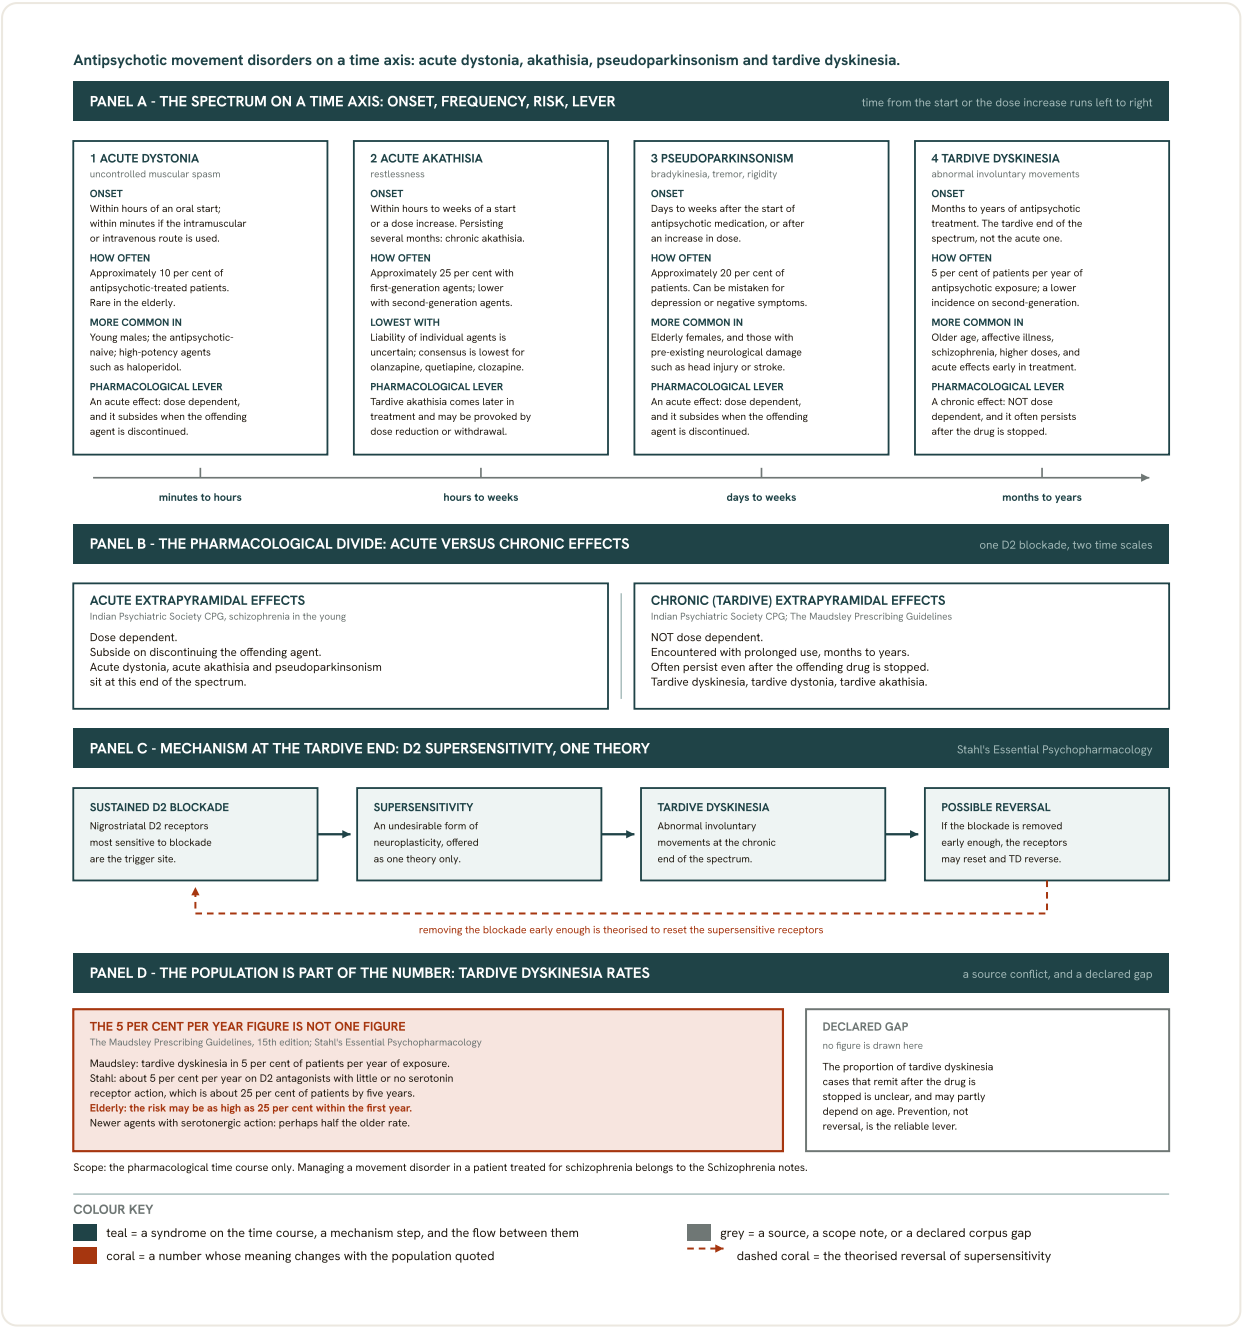

Fig 36.3 **The antipsychotic movement-disorder spectrum read on a time axis.** The four syndromes separate first by when they appear: acute dystonia within hours of an oral start and within minutes by the intramuscular or intravenous route (approximately 10 per cent of patients, more common in young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

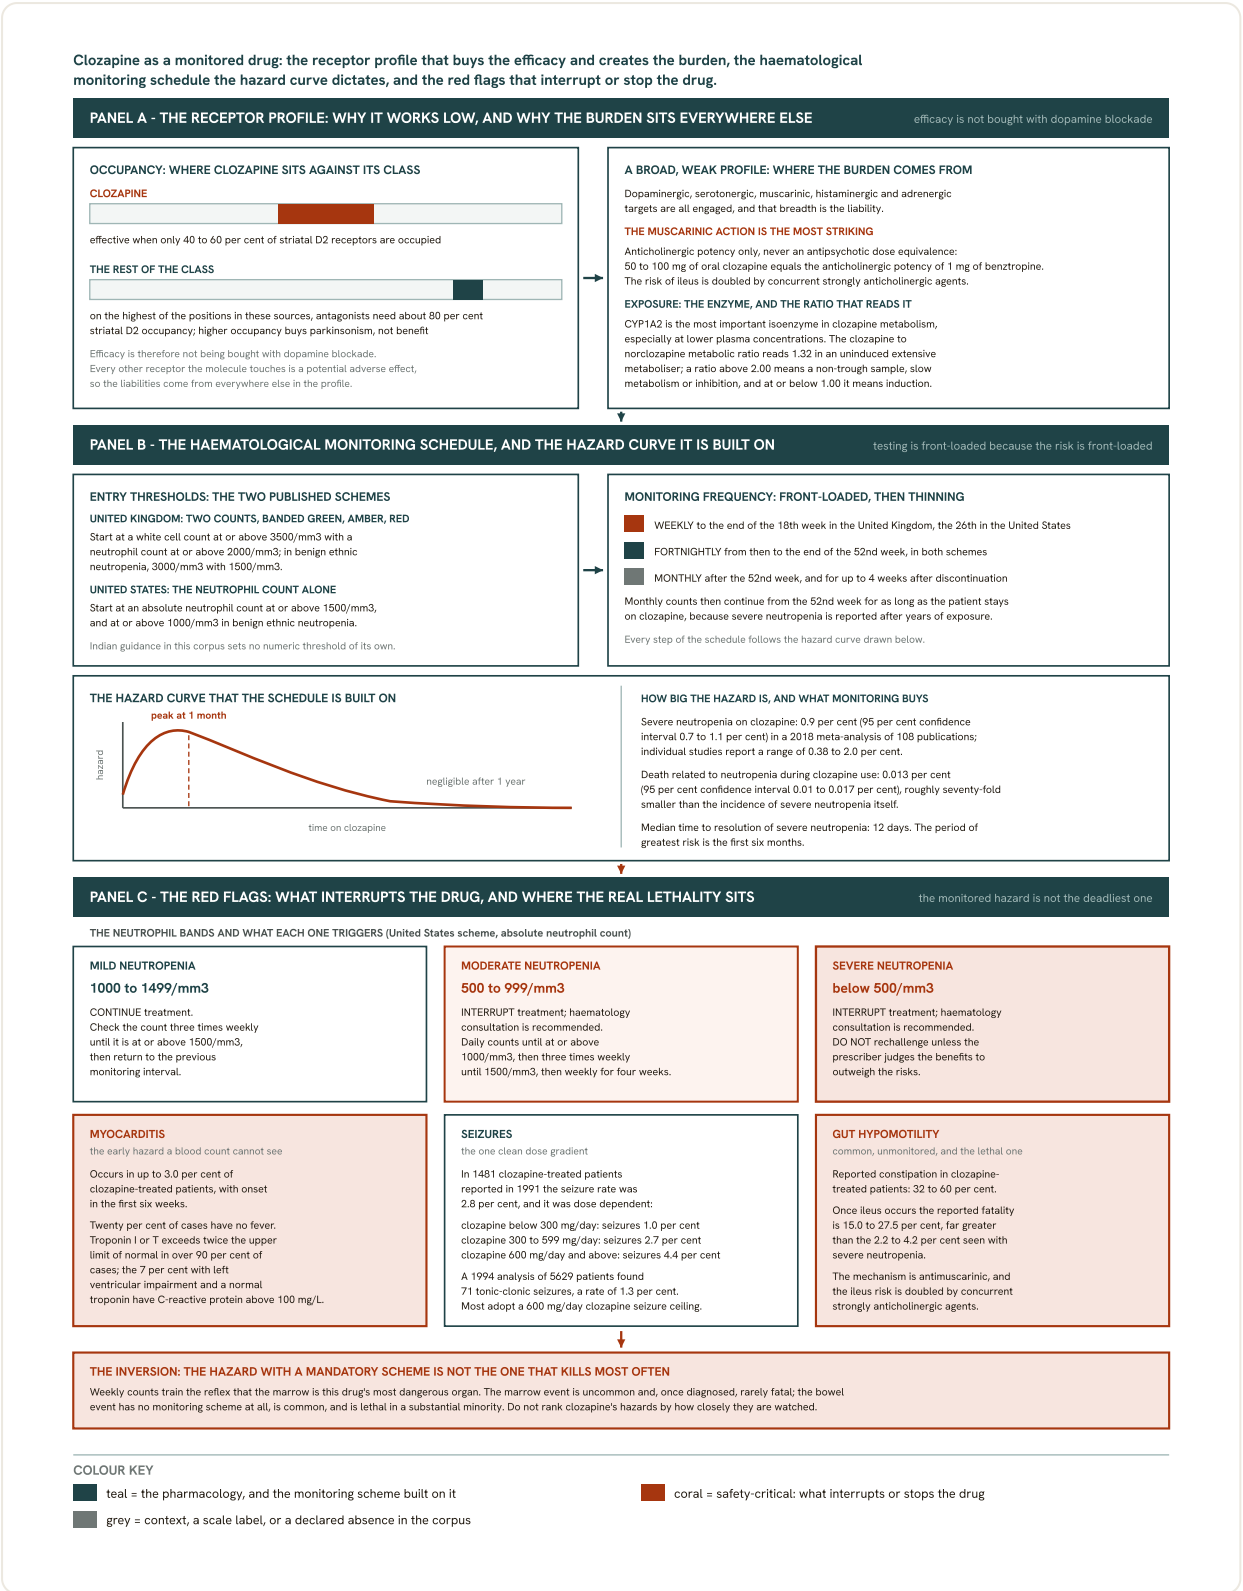

Fig 36.4 Clozapine: the profile that buys the efficacy, the monitoring schedule the hazard curve dictates, and the red flags that stop the drug. Clozapine is effective when only 40 to 60 per cent of striatal D2 receptors are occupied, where the antagonists as a group need approximately 80 per cent, so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

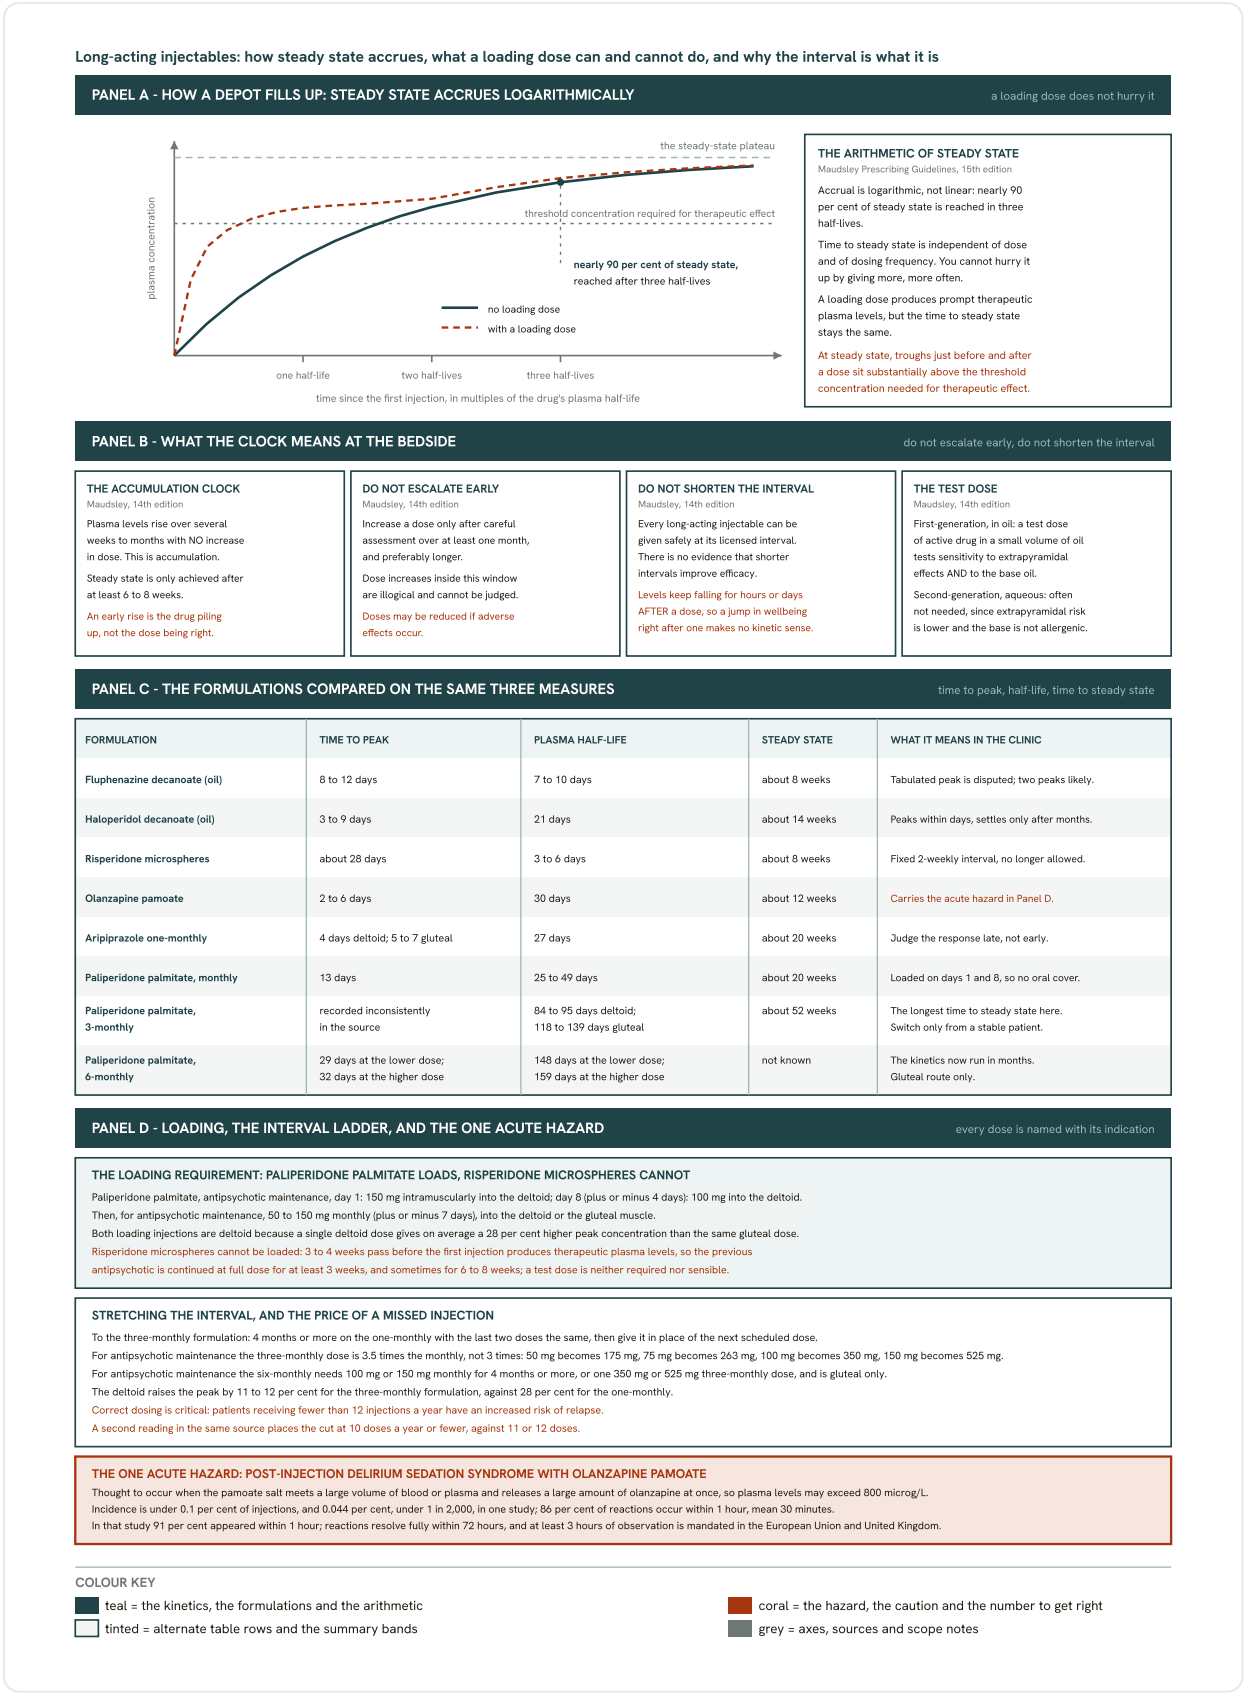

Fig 36.5 **Long-acting injectable release kinetics.** Steady state accrues logarithmically: nearly 90 per cent of it is reached in three half-lives, and the time to steady state is independent of dose and of dosing frequency, so a loading dose buys prompt therapeutic plasma levels without shortening 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.

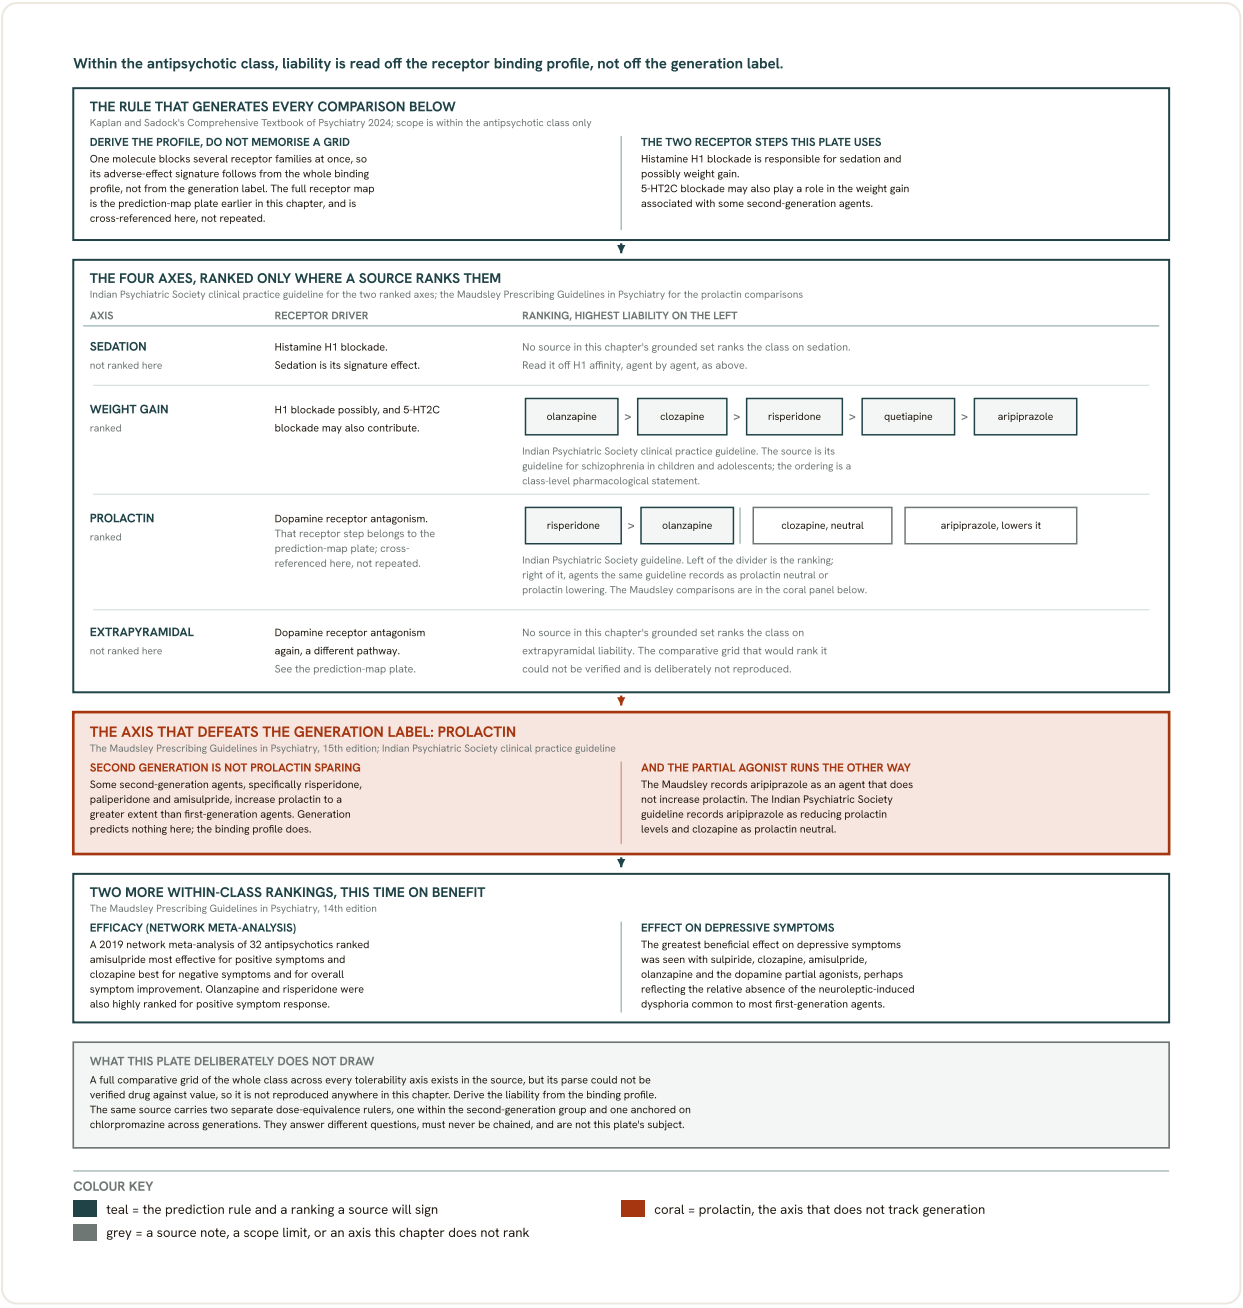

Fig 36.6 Within the antipsychotic class: what the receptor profile predicts, and where a source will actually rank the agents. Weight gain and prolactin are ranked where the Indian Psychiatric Society guideline ranks them, while sedation and extrapyramidal liability are drawn as unranked, because no 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% STRIATAL TRANSPORTER OCCUPANCY AT MINIMUM THERAPEUTIC DOSES	**Four to six days** FLUOXETINE HALF-LIFE ON CHRONIC DOSING	**4 to 16 days** ITS ACTIVE METABOLITE, WHATEVER THE DURATION OF DOSING	**Two to four days** BEFORE DISCONTINUATION SYMPTOMS APPEAR AFTER STOPPING

14.1 One target, and a threshold rather than a dose

Every drug in this class does the same primary thing: it blocks the presynaptic serotonin transporter and leaves the rest of the receptor landscape largely alone. The consequence that reframes the whole class is that the quantity which has to be achieved for an antidepressant effect is not a milligram figure at all, but a proportion of transporter occupied. Positron emission tomography performed on citalopram, fluoxetine, sertraline, paroxetine and extended-release venlafaxine, each given at its own minimum therapeutic dose, shows **about 80% occupancy** of serotonin reuptake sites in the striatum. Different molecules, different milligram figures, one shared **occupancy threshold**.

The consequences reach in opposite directions. Comparing agents in milligrams is meaningless, because a milligram figure is a property of the molecule, of its potency and of how the body handles it, and not of the target that has to be occupied; the sentence "sertraline is a weaker drug than paroxetine because the tablets are larger" is a pharmacological error rather than an approximation. Occupancy is also regional and not uniform, and that regional variation carries a teaching point of its own: occupancy in the orbitofrontal cortex runs lower than in the striatum, which is the stated explanation for obsessive-compulsive disorder requiring higher doses than depressive illness does. That is an observation about where the drug binds and how much of it binds there, not a statement about how any disorder is managed.

14.2 Receptor cleanliness, and the axes that survive it

Receptor cleanliness is the property that made this class replace its predecessor, and it is worth stating as a negative fact because that is how it behaves clinically. Most SSRIs carry essentially no clinically meaningful antagonism at muscarinic, histaminergic, adrenergic or postsynaptic

serotonin receptors, and that absence is the entire reason the class is tolerated where the tricyclics are not. The point is worth labouring: the burden that makes a tricyclic hard to take is not serotonergic at all. It is the sum of what that molecule does at every receptor it was never aimed at, and an agent that reaches the transporter without visiting those receptors inherits the benefit and not the cost. The mapping from a binding profile to the adverse effect the patient reports is set out in the section of these notes on reading a receptor profile, and the receptor promiscuity of the tricyclics themselves is taught with that class.

Cleanliness at receptors is not the same as sameness, and what survives it is the material this section exists to teach.

- **Off-target binding.** A minority of the class does visit other sites, and the sites visited predict the complaint the patient brings back.
- **Elimination kinetics.** Serum half-lives across the class run from hours to days, and the position an agent occupies on that spread predicts what stopping it will feel like.
- **Enzyme inhibition.** The agents differ in how strongly they inhibit the cytochrome P450 isoenzymes, and it is this, not the serotonergic action, that makes one SSRI a harder drug to add to an existing regimen than another. The map of the isoenzymes themselves, and of what varies in them, is set out in the section of these notes on the psychotropic substrate map; which SSRI inhibits which isoenzyme is not carried by the sources used here. The interaction matrix spanning drug classes belongs to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

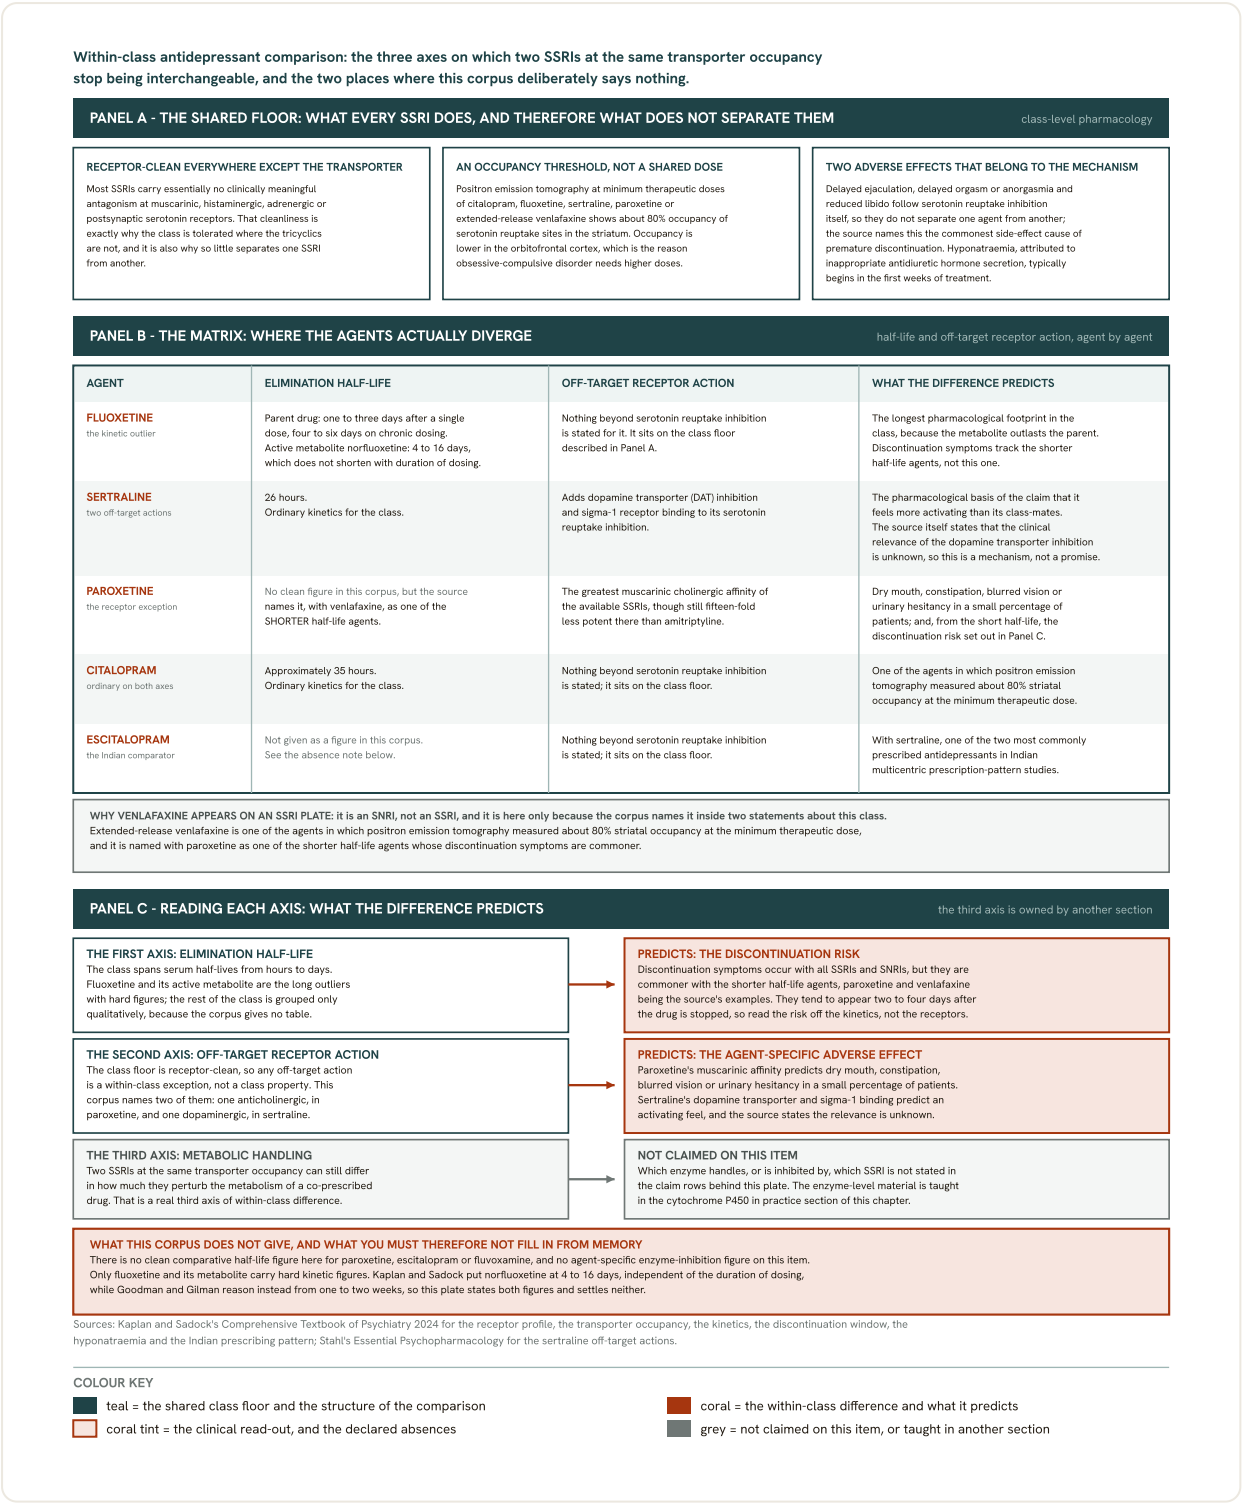

Fig 36.7 Within-class antidepressant comparison: half-life, off-target receptor action, and the axis this item does not claim. Most SSRIs are receptor-clean beyond the transporter, and every agent has to reach an occupancy threshold rather than a shared milligram dose, which is why sexual dysfunction and early hyponatraemia belong to the mechanism and not to any one drug. What separates the agents is kinetic and receptor-level: fluoxetine and its active metabolite carry the long half-lives, sertraline adds dopamine transporter and sigma-1 binding, paroxetine carries the greatest muscarinic affinity in the class, and the shorter half-life agents carry the discontinuation risk. The third axis, metabolic handling at the cytochrome P450 enzymes, is a genuine source of within-class difference that this item does not claim and that the cytochrome P450 in practice section of this chapter teaches. (Kaplan and Sadock's Comprehensive Textbook of Psychiatry 2024; Stahl's Essential Psychopharmacology.)

The accompanying figure draws the comparison out. Read the class this way and the examination question stops being a memory test and becomes a reasoning one: given an agent and a

complaint, which axis is the complaint sitting on.

RULE **What the SSRIs share is the target; what separates them is what else they touch, and how long they take to leave.** Every within-class comparison in this class reduces to that sentence. Shared mechanism generates the effects the whole class produces, and it therefore cannot be escaped by moving between agents. Off-target binding and elimination kinetics generate the differences, and those are the only differences a change of agent can buy.

14.3 The agents that break the class rule

The exceptions to that cleanliness attach to named agents, and the examiners know which agents they are. The first difference is **muscarinic cholinergic affinity**. Kaplan and Sadock's Comprehensive Textbook of Psychiatry, 11th edition records that paroxetine has the greatest affinity for muscarinic receptors of the available SSRIs, and therefore produces dry mouth, constipation, blurred vision or urinary hesitancy in a small percentage of patients. Keep that fact in proportion. The proportion of patients affected is small, so the anticholinergic profile of paroxetine is a within-class difference and not a class-changing one. And the affinity, though highest in the class, is modest in absolute terms: the same source records paroxetine as fifteen-fold less potent at muscarinic receptors than amitriptyline. Paroxetine is the anticholinergic SSRI; it is not an anticholinergic drug in the sense a tricyclic is.

The second exception runs the other way. Stahl's Essential Psychopharmacology describes sertraline as adding **dopamine transporter inhibition** and sigma-1 receptor binding to its serotonin reuptake inhibition, which is the pharmacological basis usually offered for the impression that sertraline feels more activating than its class-mates. Handle this one honestly in an answer, because the source does: the same passage states that the clinical relevance of that dopamine transporter inhibition is unknown. A mechanism that is present at the molecule and unproven at the bedside is exactly the kind of claim a viva rewards you for labelling correctly. Say that the binding is documented, say what it might explain, and say that the link to the clinical impression is not established.

MNEMONIC

Paroxetine, Sertraline, Fluoxetine

These names carry almost every within-class difference worth printing. **Paroxetine** is the one with muscarinic company, and the one whose short kinetics make stopping it uncomfortable. **Sertraline** is the one with company at another transporter. **Fluoxetine** is the kinetic outlier, the drug that is still present long after the patient has stopped taking it. Everything else in the class is, for examination purposes, the class.

14.4 Half-life: the axis with the widest spread

The **elimination half-life** is where this class separates most dramatically, and one agent accounts for most of the separation. Fluoxetine has an elimination half-life of **one to three days after a single dose and four to six days on chronic dosing**. The source gives both figures and not the reason they differ, and the reason is worth getting the right way round: a half-life lengthens when

clearance falls or the volume of distribution widens, which is the rule the section of these notes on pharmacokinetics and half-life sets out. Accumulation is what a long half-life then produces, not what lengthened it. That figure alone would make fluoxetine unusual. What makes it exceptional is the metabolite.

Norfluoxetine is pharmacologically active, and its elimination half-life is longer than the parent compound's and does not shorten however long the drug has been taken. The clinical translation is that fluoxetine's pharmacological footprint outlasts fluoxetine. Weeks after the last tablet, a patient may still be carrying a meaningful serotonergic load, which is why every question about washout, about starting a drug that interacts, and about the interval before a monoamine oxidase inhibitor, is answered differently for this agent than for the rest of the class.

The same arithmetic settles a smaller question that comes up constantly. A missed dose of a short half-life agent is a real fall in exposure, because a meaningful fraction of what is circulating is gone before the next tablet is due. A missed dose of fluoxetine is close to arithmetically invisible, because the interval between doses is a small fraction of the time the drug needs to halve. The clinical texture that follows from this, the impression that some of these drugs are forgiving and others are not, is not a property of the receptor at all.

A recital of half-life figures for the remaining agents is worth less than the principle they illustrate, and the principle is short. Serum half-lives across the class span from hours to days; fluoxetine and its active metabolite are the long outliers, and they are the ones with firm figures attached; and the gap between fluoxetine and its class-mates is several-fold rather than an order of magnitude. An examiner asking you to rank the class by half-life wants the outlier named and the direction right, not a table recited.

TRAP **Quoting a single settled half-life for norfluoxetine.** The standard references do not agree, and a candidate who has met only one of them assumes the other is a misprint and defends the wrong thing. Kaplan and Sadock's Comprehensive Textbook of Psychiatry, 11th edition puts the elimination half-life of the active metabolite at **4 to 16 days** and states that this is independent of how long the drug has been given. Goodman and Gilman's The Pharmacological Basis of Therapeutics reasons instead from one to two weeks when it sets the interval that must elapse before a monoamine oxidase inhibitor may be started. Neither source is wrong and the ranges overlap: the wider figure admits more between-individual variation than the working number the washout arithmetic is built on. The safe answer gives the range, names whose figure it is, and does not manufacture agreement. The washout interval itself is taught with the monoamine oxidase inhibitors.

14.5 Discontinuation, read off the kinetics

Discontinuation symptoms are the clearest illustration in the whole class of a clinical phenomenon that is predicted by pharmacokinetics and by nothing else. They can occur with all SSRIs and SNRIs, but are more common with the shorter half-life agents, paroxetine and venlafaxine being the named examples. Nothing about a receptor profile predicts which agents will be difficult to stop; the elimination curve predicts it completely. The agent at the other end of

the spread is fluoxetine, whose long tail means the concentration falls slowly enough to taper the patient without anyone intending it.

The mechanism underneath is worth stating because it makes the rule reconstructible rather than memorised. Continuous occupancy of the transporter over weeks of treatment is met by adaptation in the systems downstream of it. When a rapidly cleared agent is stopped, occupancy falls faster than those adaptations can unwind, and the mismatch is what the patient experiences. When a slowly cleared agent is stopped, occupancy falls under cover of the same adaptations reversing, and there is no mismatch to feel. The half-life is therefore not a proxy for the risk; it is the risk, expressed as a rate.

The timing carries the same lesson. Symptoms tend to appear **two to four days** after the medication is stopped, which is an interval measured in the drug's own elimination window. That is the discriminator worth carrying: a relapse of the underlying illness does not respect the pharmacokinetics of the tablet, and a syndrome that arrives inside the clearance window and resolves when the drug is reinstated is a kinetic event. The clinical handling of that event in a patient being treated for a depressive episode belongs to the Mood disorders: pharmacotherapy notes; what belongs here is the reason the risk differs between agents at all.

14.6 The effects that belong to the mechanism, not the molecule

Some adverse effects of this class are not properties of any agent within it. They are properties of inhibiting serotonin reuptake, and they follow the mechanism wherever it goes. **Sexual dysfunction** is the important one. Delayed ejaculation, delayed orgasm or anorgasmia and reduced libido are a common effect of serotonin reuptake inhibition itself, which is why the same complaint recurs across agents that share nothing else except the target. It is also the commonest side-effect reason for premature discontinuation, which is the fact that makes this more than a footnote: the effect that most often ends treatment is the one least amenable to changing the tablet.

Hyponatraemia behaves as a class hazard in the same way. SSRIs raise the risk of hyponatraemia, and it typically begins during the first weeks of treatment; the source offers inappropriate antidiuretic hormone secretion as the probable mechanism rather than a settled one, and an answer that hardens it into the mechanism has said more than the text does. Of the pair, the timing is both the firmer statement and the more useful one, because it tells you which stretch of treatment the explanation is available for. The monitoring schedule that follows from this across the psychotropics belongs to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

PITFALL

Reading sexual dysfunction as a property of the particular molecule. It is a property of serotonin reuptake inhibition, so the pharmacological expectation is that a move to another agent doing the same thing at the same transporter carries the same liability. Clinicians who have not held that distinction attribute the failure of the switch to the patient, or to the illness, rather than to the mechanism they never left.

GOOD TO REMEMBER

Because the sodium effect begins early, new confusion, unsteadiness or an unexplained fall arising in the first weeks after an SSRI is started has a pharmacological explanation sitting immediately available. Later in treatment the same presentation still deserves a sodium, but the drug is a weaker candidate for it, and reaching for the prescription first can cost the real diagnosis.

14.7 Class property or agent property: the question to ask of every finding

One question resolves most of this material under examination conditions. Given an adverse effect, a kinetic behaviour or a complaint, ask whether it belongs to serotonin reuptake inhibition or to the particular molecule doing it. The answer decides whether moving within the class can help at all, and it is the single most reliable way to reconstruct the comparison when the details have faded.

PROPERTY	CLASS OR AGENT	WHAT THE STANDARD TEXTS ANCHOR
The quantity that must be reached for effect	Class	A shared occupancy threshold rather than a shared milligram dose
Sexual dysfunction	Class	Follows serotonin reuptake inhibition itself, not the individual agent
Hyponatraemia	Class	Begins early in treatment; inappropriate antidiuretic hormone secretion offered as the probable mechanism
Muscarinic adverse effects	Agent	Paroxetine carries the greatest muscarinic affinity of the available agents
Off-target transporter and sigma binding	Agent	Sertraline adds dopamine transporter inhibition and sigma-1 receptor binding
Elimination half-life	Agent	Fluoxetine and its active metabolite are the long outliers
Discontinuation liability	Agent, read off the kinetics	More common with the shorter half-life agents

IN INDIA

Kaplan and Sadock's Comprehensive Textbook of Psychiatry, 11th edition, India chapter reports that multicentric Indian prescription-pattern studies find escitalopram and sertraline to be the most commonly prescribed antidepressants, which makes that pair the within-class comparison worth rehearsing before any other. It changes what you do with the material above. Neither of those agents is the muscarinic outlier, so a patient of yours on one of them who reports dry mouth, constipation or urinary hesitancy should send you looking for another cause rather than accepting the anticholinergic explanation that fits paroxetine. Neither is the kinetic outlier either, so the slow self-taper that fluoxetine supplies without being asked is not available for the pair you will most often be stopping, and the discontinuation question has to be raised deliberately rather than assumed away.

Sexual dysfunction and hyponatraemia belong to the mechanism; muscarinic burden, off-target binding, enzyme inhibition and half-life belong to the molecule.

15 · Beyond the SSRIs: SNRI, mirtazapine, bupropion, vortioxetine and agomelatine

Everything gathered here is defined by what it does instead of, or in addition to, blocking a serotonin transporter. The serotonin reuptake inhibitors are one pharmacological idea executed with small variations. The agents in this section are not: they share only a negative property, that none of them is a plain serotonin reuptake inhibitor, and past that point they diverge at the level of the molecule. One group inhibits two transporters rather than one. One raises monoamine transmission without touching a transporter at all. One works outside the serotonin system altogether. One adds direct receptor actions to a transporter block. One leaves the monoamine reuptake axis entirely. Examinations reward the candidate who can state, for any single agent here, the one molecular action from which the rest of its behaviour follows.

Two habits of thought are worth breaking before starting. The first is reading a class label as a description of a molecule: an abbreviation names the axis a drug was designed around, not the target it actually engages hardest, and the drugs inside one label do not engage their two transporters in the same order or to the same degree. The second is treating mechanism as a fixed property of a drug rather than of a drug at a dose, which for at least one agent in this group is simply false. Where each of these medicines belongs in the treatment of a depressive episode is the business of the Mood disorders: pharmacotherapy notes. What follows is the pharmacology that makes those decisions intelligible.

Two transporters	alpha-2	NA and DA	Melatonergic
THE SNRI AXIS, ENGAGED IN A DIFFERENT ORDER BY DIFFERENT AGENTS	MIRTAZAPINE: RELEASE RAISED, NO TRANSPORTER BLOCKED	BUPROPION, OUTSIDE THE SEROTONERGIC GROUP	AGOMELATINE, OFF THE REUPTAKE AXIS ENTIRELY

15.1 The SNRI label and the split it conceals

The abbreviation promises a drug that inhibits two monoamine transporters, and each member does, but it says nothing about which transporter the molecule reaches first or holds most tightly.

That property, the drug's **transporter preference**, is what separates the members from one another. Venlafaxine, duloxetine and desvenlafaxine all reach for the serotonin transporter first; levomilnacipran is the exception, working predominantly at the noradrenaline transporter. The label therefore conceals a genuine pharmacological split rather than a difference in potency, and a candidate who recites the class as though it were one object has already lost the discriminating detail.

The consequence is worth stating plainly. A serotonin-preferring member, at the doses in ordinary use, is doing most of its work through the same transporter an SSRI works through, with a noradrenergic component layered on top. A noradrenaline-preferring member leads with noradrenergic pharmacology and picks up serotonergic action second. Two drugs bearing one abbreviation therefore start their effect, and start their adverse effects, from opposite ends of the same axis. Which of them suits which patient is a treatment decision and belongs to the Mood disorders: pharmacotherapy notes; the pharmacological point owned here is only that the label is not the mechanism.

RULE **A class name states the axis a drug was designed around; it does not state which target the molecule engages, or at what dose it engages it.** Inside one abbreviation, one member leads with noradrenaline while the rest lead with serotonin, and at least one member changes which transporter it is meaningfully inhibiting as the dose climbs. Read the molecule, then the dose, then the label, and never in the reverse order.

15.2 Dose as a mechanism switch: venlafaxine against duloxetine

Venlafaxine is the reason mechanism has to be quoted with a dose attached. Its **dose-dependent dual action** means the serotonergic effect is present throughout the range while the noradrenergic effect is not: noradrenaline reuptake inhibition is seen only when the antidepressant dose is higher than **150 mg per day**, a threshold recorded here as mechanism because the prescribing that acts on it is owned by the Mood disorders: pharmacotherapy notes, so a prescription below that level is pharmacologically a serotonin reuptake inhibitor carrying an SNRI's name. Nothing about the drug's identity has changed; what has changed is which of its two possible actions is actually being engaged in that patient.

Duloxetine is the within-class contrast that makes the point sharp. Both transporters are inhibited from its starting doses, so the second mechanism is available without having to titrate towards it. That is a difference in kind, not a difference in potency: the two drugs do not sit at different points on one curve, they have different relationships between dose and mechanism. Stated as an examination sentence, venlafaxine acquires its second transporter with dose and duloxetine has it from the outset.

Below 150 mg per day the venlafaxine antidepressant prescription is pharmacologically a serotonin reuptake inhibitor wearing an SNRI label, a statement about mechanism and not a prescribing instruction, since prescribing is not taught here; duloxetine is dual from its starting dose.

TRAP **Answering that a patient on a low dose of venlafaxine is receiving dual-action treatment.** The error is easy to make because the class name is on the packet at every strength, and because mechanism is memorised as a property of drugs rather than of a drug at a dose. Once the threshold is crossed the dual description is accurate; below it, every line of reasoning that depends on noradrenergic action fails, including any expectation of a different adverse-effect profile from that of an SSRI. The same trap runs in reverse in a viva: asked how to obtain noradrenergic action, the answer is not to add a second drug before the first has been taken past its own threshold.

15.3 Mirtazapine: more monoamine without a transporter

Mirtazapine is the cleanest demonstration in psychiatry that raising a transmitter does not require blocking its reuptake. The mechanism begins at a presynaptic inhibitory receptor: blockade of alpha-2 noradrenergic receptors raises the release of serotonin and of noradrenaline together. Those receptors restrain transmitter output, so an antagonist there lifts a brake and release rises. Synaptic transmitter concentration climbs, as it does under a reuptake inhibitor, but by an entirely different route, and the designation the reference literature gives this drug, a **noradrenergic and specific serotonergic antidepressant**, encodes precisely that difference.

Why this distinction earns marks rather than merely being tidy: a mechanism that works by disinhibiting release behaves differently from one that works by obstructing clearance. Reuptake blockade acts wherever the transporter is expressed and depends on the transporter being there; release enhancement depends on there being tonic inhibition to lift, and therefore on the firing of the neurone itself. The two mechanisms also combine rather than duplicate, which is the pharmacological reason a release-enhancing agent and a reuptake inhibitor are not simply two versions of the same intervention. The reasoning behind combining agents in a resistant depressive illness sits with the Mood disorders: pharmacotherapy notes.

15.4 The same drug's adverse-effect signature, read off the same list

The alpha-2 blockade is only half of this molecule's pharmacology, and the other half is where the patient's complaints come from. The same compound is a postsynaptic antagonist at 5-HT2, 5-HT3 and histaminergic receptors, with a weaker action at alpha-1 adrenergic and muscarinic receptors, and the reference literature lays the drug's characteristic adverse effects at the door of that second set rather than of the presynaptic action described above. The general method of reading a receptor profile forward into a predicted adverse-effect signature is taught elsewhere in these notes; what belongs here is that this drug is its cleanest worked example, because therapeutic mechanism and adverse-effect signature sit in two separable halves of one binding profile.

- Sedation, attributed to blockade of histaminergic receptors, and the reason this agent is described as the sedating member of the group.
- Postural symptoms, from the milder antagonism at alpha-1 adrenergic receptors.
- Dry mouth and the rest of the anticholinergic complaints, from the milder muscarinic antagonism.

- The comparative absence of the adverse-effect pattern a pure serotonin reuptake inhibitor produces, attributed to blockade of 5-HT2 and 5-HT3 receptors, which prevents unopposed stimulation of those receptor populations.

Read as a whole, the profile explains a drug that raises two monoamines and yet does not behave like the drugs that raise them by transporter blockade. That is the entire teaching point, and it is reconstructible from the receptor list without memorising an adverse-effect table.

15.5 Bupropion: the antidepressant that is not serotonergic

One member of this group leaves serotonin alone. Bupropion's effects are predominantly noradrenergic and dopaminergic, and the drug is activating rather than sedating, which is the reason it is so often prescribed alongside a more sedating serotonin-enhancing antidepressant. Calling it an **activating antidepressant** is a statement about the transmitter systems engaged rather than about mood: dopaminergic and noradrenergic action drives arousal, drive and initiation, and an agent that raises those without raising serotonin has a profile unlike anything else in the antidepressant sections of this chapter. The complementary pairing follows from that fact rather than from empirical habit.

The liability that defines the molecule is seizure risk, and it is instructive precisely because it is not a fixed property of the drug. The seizure risk is dose related and is greatest with instant-release formulations, and it is lower with **slow-release** formulations at antidepressant doses **under 300 mg per day**, as stated in The Maudsley Prescribing Guidelines. Two variables therefore govern it, the amount given and the rate at which the preparation delivers it, which is a formulation-kinetics statement rather than a statement about the compound. Recognition and acute management of a drug-induced seizure belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

> **GOOD TO REMEMBER**
>
> When a seizure risk is quoted for this drug, ask which preparation the figure refers to before using it. Instant-release and slow-release preparations do not carry the same risk at the same total daily amount, so a number carried across from one formulation to the other is not conservative, it is simply wrong in whichever direction the formulations differ. The same discipline applies to any drug whose hazard is driven by peak concentration rather than by total exposure.

15.6 Vortioxetine and agomelatine: modulation past the transporter

The newest members of the group are the ones whose names carry the least information, so their mechanisms have to be learned directly. Vortioxetine is described as **multimodal**, and the word carries a precise pharmacological content: serotonin transporter blockade sits alongside antagonism at 5-HT3 and 5-HT7 and agonism at 5-HT1A and 5-HT1B. One transmitter system is therefore modulated at several points at once rather than only at its transporter. Whether that translates into a distinguishable clinical advantage is a separate question, and one this section deliberately does not answer; the pharmacology is what is owned here.

Agomelatine goes further and steps off the axis altogether. It is a **melatonin receptor agonist** and, at the same time, an antagonist at the 5-HT2C receptor, which places it outside the reuptake and the enzyme families entirely. Alone in this section, its primary target is neither a monoamine transporter nor a monoamine receptor but a melatonin receptor, and mirtazapine is the instructive contrast: that drug also raises transmission without a transporter, but it does so at a monoamine receptor, so it never leaves the monoamine system. That structural oddity is the examinable fact, and it also explains why the drug's cautions do not resemble those of its neighbours: what constrains its use is hepatic, not serotonergic. The specific liver-function testing schedule, which tests are taken, at what intervals, and the transaminase level at which the drug is stopped, is not taught here; monitoring across the psychotropics belongs to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes, and a schedule quoted from memory is exactly the kind of number that reaches a patient wrongly.

15.7 Trazodone, and dose-effect separation inside one molecule

Trazodone earns its place here for a reason that is pharmacological rather than therapeutic: it is the standard illustration that one molecule can have two dose ranges with two different purposes. Its sedative load is heavy enough that the drug is far more often written up at **50 to 100 mg**, a range the standard reference text explicitly calls subtherapeutic for antidepressant purposes, in order to treat insomnia rather than the depressive illness itself, a use whose prescribing sits with the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes. The same compound at a much higher exposure is an antidepressant; at this exposure it is a hypnotic, and calling it an antidepressant at either dose without saying which is being used is imprecise to the point of being misleading.

Generalised, the principle is that a **subtherapeutic antidepressant dose** is not a small version of a therapeutic one. It is a different pharmacological intervention, engaging whichever actions have the lowest threshold, in this case the sedating ones, and leaving the actions that require higher exposure unengaged. That is the same logic met earlier with venlafaxine, seen from the other side: there, the higher-threshold action was the one the class name advertised; here it is the one the drug is nominally licensed for.

> **PITFALL**
>
> Counting a hypnotic-dose trazodone prescription as an antidepressant trial. It appears in the medication history as an antidepressant, is often written in the evening alongside another agent, and is then carried forward into a treatment-resistance calculation as though the illness had failed to respond to it. It has not been tried at an antidepressant exposure at all, and the same reasoning error recurs with any drug whose low-dose and full-dose indications differ.

15.8 The whole set on one axis

Laid side by side, these agents answer a single question in different ways: what does the molecule do that a serotonin reuptake inhibitor does not. The accompanying comparison is best read down the middle column, because the defining action in each row predicts the entries either side of it.

AGENT	THE ACTION THAT DEFINES IT	WHAT FOLLOWS FROM THAT ACTION
Venlafaxine	Serotonin transporter first, with noradrenaline reuptake inhibition only above the threshold dose	Mechanism moves with dose, so the class name describes only the upper part of the range
Duloxetine	Both transporters inhibited from the starting dose	Dual action without titrating towards it; the within-class contrast to venlafaxine
Desvenlafaxine	Serotonin transporter preferred	Sits with venlafaxine on the serotonin-leading side of the class
Levomilnacipran	Noradrenaline transporter predominantly blocked	The one member that leads with noradrenaline rather than serotonin
Mirtazapine	Blockade of alpha-2 receptors raising release, plus postsynaptic 5-HT2, 5-HT3, histaminergic, alpha-1 and muscarinic antagonism	Monoamines rise with no transporter blocked; the postsynaptic list is the adverse-effect list
Bupropion	Predominantly noradrenergic and dopaminergic	Outside the serotonergic agents; activating rather than sedating, and paired accordingly
Vortioxetine	Serotonin transporter blockade plus 5-HT3 and 5-HT7 antagonism with 5-HT1A and 5-HT1B agonism	Modulation at several points in one transmitter system, not at the transporter alone
Agomelatine	Melatonin receptor agonism with 5-HT2C antagonism	Outside the reuptake and the enzyme families entirely; its cautions are hepatic
Trazodone	A sedative load heavy enough to separate its hypnotic use from its antidepressant use	Usually encountered well below an antidepressant exposure

MNEMONIC

Release, Activate, Multimodal, Melatonergic

The exits from the reuptake axis, in the order this section takes them. **R**elease: mirtazapine raises serotonin and noradrenaline by blocking an inhibitory adrenergic receptor. **A**ctivate: bupropion works through noradrenergic and dopaminergic effects and behaves as an activating agent. **M**ultimodal: vortioxetine adds direct receptor agonism and antagonism to a transporter block. **M**elatonergic: agomelatine leaves the monoamine reuptake families altogether.

IN INDIA

Not every agent in this section is one a trainee will actually put on a prescription. Agomelatine, vilazodone, vortioxetine and levomilnacipran are barely used in Indian practice, a pattern the Indian source literature attributes to cost or to efficacy that did not match expectation. The working consequence is where to place effort: the within-class choices that recur on Indian wards fall among the serotonin-preferring SNRIs, mirtazapine and bupropion, and those are the profiles to hold at recall depth, including the venlafaxine threshold and the bupropion formulation rule. The scarcer molecules remain fully examinable and are learned here as pharmacology rather than met at the bedside, so unfamiliarity with a drug on the ward is not a reason to leave its mechanism unlearned.

The unifying discipline, finally, is to insist on a mechanism sentence for every one of these agents before anything else is said about it. Class membership, sedation, activation and tolerability are all downstream of the molecular action, and every discriminating point an examiner can build, including the dose threshold, the transporter split, the release-versus-reuptake distinction and the formulation-dependent hazard, is recoverable from that one sentence. Cross-class comparison, interaction burden and prescribing in the medically ill are addressed in the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

16 · Tricyclics: why the class is dangerous and when it still wins

The tricyclics are the class that teaches receptor promiscuity, because almost everything that makes them dangerous is something the molecule does that has nothing to do with treating anything. A tricyclic blocks monoamine reuptake, and if that were all it did the class would look much like the selective agents that displaced it. It is not all it does. The same molecule sits on muscarinic, histaminergic and adrenergic receptors, and on a voltage-gated ion channel the heart and the brain both depend on, and those additional actions are the whole story of why a family of drugs with no efficacy problem was pushed to the back of the formulary.

The marks here are rarely for reciting a tricyclic's adverse effects. They are for the derivations underneath them: which structural feature predicts which monoamine, why prescribing one tricyclic exposes the patient to a pair of pharmacologies, which unwanted action actually kills, and why an identical dose produces very different plasma concentrations in patients of the same weight. Whether a tricyclic is the right agent for a particular patient with a depressive episode belongs to the Mood disorders: pharmacotherapy notes; overdose management, the plasma target ranges and prescribing in established cardiac disease belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

4	100 ms	40-fold	30-day
UNWANTED ACTIONS EVERY MEMBER CARRIES BEYOND REUPTAKE BLOCK	QRS WIDTH ABOVE WHICH OVERDOSE SEIZURES ARE PREDICTED	SPREAD IN PLASMA CONCENTRATION ON IDENTICAL DOSING	SUPPLY IS APPROXIMATELY THE LETHAL DOSE

16.1 The structural split that predicts which monoamine

The class divides on a single feature of the side chain, and that feature is not cosmetic: it predicts which transporter the molecule prefers. Kaplan and Sadock's Comprehensive Textbook of Psychiatry states the mapping directly. The **tertiary amines** have the greater affinity for the **serotonin transporter**, whereas the **secondary amines** primarily block the uptake of noradrenaline, so the degree of methylation on the side chain is what predicts which monoamine the drug engages.

That converts an unmanageable list into a derivation. Memorise which agent is which and the first stem naming an unrevised compound defeats you; know that the fuller methylation goes with the serotonergic profile and any compound can be placed from its structure. The split also predicts the metabolic route, so a patient's enzyme phenotype does not affect every member equally.

AXIS OF COMPARISON	TERTIARY AMINES	SECONDARY AMINES
Side-chain nitrogen	The more fully methylated form	The demethylated product of the corresponding tertiary compound
Transporter preferred	Greater affinity for the serotonin transporter	Primarily blocks the uptake of noradrenaline
Worked examples	Amitriptyline and imipramine	Nortriptyline and desipramine respectively
Standing in the body	Parent compound, converted in vivo	Also circulates as the active metabolite of its parent
Enzyme whose deficiency bites hardest	CYP2C19, which metabolises the tertiary compounds	CYP2D6, which hydroxylates desipramine and nortriptyline

16.2 Demethylation: prescribing a parent and its metabolite together

The structural split would be a tidy dichotomy if the halves of the class stayed separate in the body. They do not. **Demethylation** converts amitriptyline to nortriptyline and imipramine to desipramine, so a patient given a tertiary amine is exposed to a serotonergic parent compound and a noradrenergic **active metabolite** at the same time. The prescription names one drug; the patient receives a mixture.

Three consequences follow. The receptor pharmacology the patient experiences is the sum of parent and metabolite rather than the profile of the named compound, which is part of why the tertiary agents feel heaviest to take. The balance of that mixture is set by enzyme activity, so an inherited metabolic difference does not simply raise or lower a level, it shifts the ratio of serotonergic to noradrenergic compound. And a secondary amine is a different proposition from its parent, because a patient given the metabolite alone never receives the parent.

TRAP **Treating the tertiary and secondary amine split as a clean division of labour, so that a tertiary agent is answered as a purely serotonergic drug.** Candidates learn the mapping between structure and transporter, which is correct, and then apply it to a patient rather than to a molecule. In the patient the mapping is blurred by metabolism: the tertiary compound is being converted to the secondary one and both are circulating. A stem naming a tertiary agent and describing a noradrenergic effect is not contradicting the structural rule; it is testing whether the rule is remembered as a statement about molecules while the patient contains a mixture.

GOOD TO REMEMBER

Moving a patient from a tertiary amine to the corresponding secondary amine is not a switch to an unrelated drug. Pharmacologically it is the withdrawal of the serotonergic parent from a pair the patient was already receiving, leaving the noradrenergic compound their own metabolism had been making all along. The expectation follows: the noradrenergic action continues, the serotonergic contribution goes, and the four shared class actions stay, because the metabolite carries them too. The burden falls in size, not in kind. Whether the move is the right treatment decision for a depressive episode belongs to the Mood disorders: pharmacotherapy notes.

16.3 The unwanted actions that define the class burden

Everything so far concerns what the class is for. The rest concerns what it costs, and the cost is a set of actions contributing nothing to the therapeutic effect. Stahl's Essential Psychopharmacology puts the accusation plainly: every drug in the class shares at least four unwanted pharmacological actions in addition to reuptake inhibition, namely blockade of **muscarinic cholinergic receptors**, of H1 histamine receptors, of **alpha-1 adrenergic receptors**, and of **voltage-sensitive sodium channels**. The same source is explicit that the limitation of this class has never been its efficacy, which is the sentence that makes the topic coherent.

Separate burden from benefit that way and the adverse-effect signature stops needing to be memorised. Muscarinic blockade withdraws parasympathetic tone from the salivary glands, ciliary muscle, bladder and gut and reduces central cholinergic transmission, which is why the complaints cluster rather than arriving as unrelated symptoms. Histamine blockade blunts central arousal and disturbs appetite regulation, so sedation and weight gain travel together. Alpha-adrenergic blockade removes the vasoconstrictor response supporting blood pressure on standing, so the dizziness is postural. The general method of converting a receptor-binding profile into the complaint a patient reports is developed elsewhere in these notes; the tricyclics are simply the class in which it is least optional. The fourth action breaks the pattern, because it is not a receptor at all.

ACTION BLOCKED	PHARMACOLOGICAL CONSEQUENCE	WHAT THE PATIENT REPORTS
Muscarinic cholinergic receptors	Parasympathetic tone withdrawn from glands, ciliary muscle, bladder and gut	Dry mouth, blurred near vision, constipation, urinary hesitancy, forgetfulness
H1 histamine receptors	Central histaminergic arousal blunted, appetite regulation disturbed	Sedation, morning heaviness, increased appetite and weight gain
Alpha-1 adrenergic receptors	Vasoconstrictor support of blood pressure on standing removed	Postural light-headedness and falls in frail patients
Voltage-sensitive sodium channels	Fast inward current slowed in cardiac conducting tissue and in neurons	Little at therapeutic doses, and the lethal features in overdose

RULE **A tricyclic treats through reuptake inhibition; everything that limits the class comes from actions that are not reuptake inhibition, and one of those actions is what kills.** Efficacy and burden are carried by the same molecule through different pharmacology, which is why the later history of the antidepressants reads as an attempt to keep the reuptake block and discard the rest. Put that way, the examinable question answers itself: why a class with no efficacy problem was displaced anyway.

MNEMONIC

MASH

The unwanted actions in one word. **M**uscarinic blockade, for the dry mouth, blurred vision, constipation and confusion. **A**lpha-adrenergic blockade, for the postural drop. **S**odium channel blockade, for the conduction effect and the overdose syndrome. **H**istaminergic blockade, for sedation and weight gain. S is the odd letter: not a receptor, and the one that kills. It delivers the whole adverse-effect signature and asserts nothing beyond the source's own list.

16.4 Membrane stabilisation: the class behaves as an antiarrhythmic

The important point about the cardiac action is not that it exists but that it has a name. The Companion to Psychiatric Studies classifies it rather than describing it. Tricyclics have the properties of **Class I antiarrhythmics**, promoting **membrane stabilisation** through inhibition of fast sodium channels, and they display both Type IA quinidine-like effects, with reduced Purkinje fibre action potential amplitude, reduced membrane responsiveness and slowed conduction, and Type IB lidocaine-like effects on action potential duration and the effective refractory period.

Naming the class is not pedantry. It says this is a recognised electrophysiological action with a characterised profile, so its consequences can be predicted rather than catalogued: slowed conduction through the ventricular conducting system is what a Class I agent does, which is why the electrocardiographic change of interest here is conduction time. It also locates the pharmacodynamic danger, since adding a drug with Class I properties to another agent with Class I properties is an interaction rather than a coincidence, and the same source warns against

that combination. Interactions across drug classes, and prescribing in established cardiac disease, belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes. What belongs here is that a psychotropic sits inside a cardiological drug classification.

16.5 One channel, heart and brain, and the marker on the tracing

The sodium-channel action is present all the time, not only after an overdose, and the sources are careful about the distinction. Tricyclics block voltage-sensitive sodium channels weakly in the heart and in the brain at therapeutic doses, and in overdose that single action is held to account for coma and seizures through central nervous system effects and for arrhythmia, cardiac arrest and death through peripheral cardiac effects. One mechanism, and no separate explanation needed for the neurological features.

Weak blockade at therapeutic concentrations is tolerated because enough current remains for normal conduction and normal neuronal excitability; as concentration rises the blockade deepens in both organs, and the deepening in cardiac conducting tissue becomes measurable on the tracing before the seizures and arrhythmias it predicts.

The bedside expression of that mechanism is a measurement. Kaplan and Sadock's Comprehensive Textbook of Psychiatry gives the threshold: in overdose, a QRS interval **longer than 100 milliseconds** may predict the development of seizures and ventricular arrhythmias. The number is the mechanism made visible: the QRS complex represents ventricular depolarisation, depolarisation is the fast sodium current, and a widened QRS is a direct readout of how much channel is blocked. What follows from crossing that threshold, including the therapy it indicates, is overdose management and belongs to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

16.6 The narrow margin, and why the prescription is the hazard

The class-limiting statement in the literature is not about adverse effects at all. Goodman and Gilman's The Pharmacological Basis of Therapeutics gives the conduction effect the decisive role: the quinidine-like effect on cardiac conduction can be life threatening in overdose and is what limits use of the class in patients with heart disease, and it is the primary reason that only a limited supply should be available to the patient at any one time. Note the structure of that sentence, because it is the structure of the topic: the effect that is tolerable at therapeutic doses is the same effect that is lethal above them.

What makes that practical rather than theoretical is arithmetic. Stahl's Essential Psychopharmacology puts it bluntly: the lethal dose of a tricyclic is only about **a 30-day supply** of the drug. A routine monthly prescription is therefore, in itself, approximately a lethal quantity, which is more uncomfortable than the familiar observation that these drugs are dangerous in overdose. Nothing unusual has to happen, and that, rather than any failure of effectiveness, is why the class fell out of first-line use.

TRAP **Attributing death in tricyclic overdose to the anticholinergic burden.** The reasoning feels sound, because the anticholinergic actions dominate the clinical picture and are usually learned first. The sources place the lethal mechanism elsewhere: blockade of voltage-sensitive sodium channels is what is held to produce the coma, the seizures, the arrhythmia and the cardiac arrest, and the conduction effect is what is named as the reason the class is limited. Anticholinergic toxicity as a syndrome, and its management, belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes. The anticholinergic actions make the presentation recognisable; the sodium channel makes the overdose fatal.

16.7 Why the same dose is not the same drug

Patients of the same weight given the same tricyclic at the same dose can end up in entirely different places, and the reason is genetic rather than behavioural. Kaplan and Sadock's Comprehensive Textbook of Psychiatry sets out the pharmacogenetics for this class specifically. Roughly **5% to 10%** of Caucasians are homozygous for the recessive CYP2D6 trait and hydroxylate desipramine and nortriptyline poorly, so a standard dose behaves in them as a much larger one.

That textbook is not internally consistent about the figure, and the discrepancy is worth carrying rather than resolving. Its tricyclic chapter gives the range quoted above; its chapter on the selective serotonin reuptake inhibitors gives a narrower band, six to eight percent, for deficiency of the same enzyme in the same population. The ranges overlap and neither refutes the other, so the honest answer names this chapter's range and acknowledges the narrower one elsewhere in the same source. Other reference tables run wider still, as the cytochrome P450 section of these notes records.

The population distribution is not the same everywhere, which is where this becomes locally load-bearing. About 50% of those of Asian descent carry the **CYP2D6*10 allele**, which reduces clearance by that enzyme and makes them intermediate metabolisers. The enzyme handling the tertiary compounds is affected separately.

All of that variability arrives in the plasma, and its magnitude is the number to remember. The variability in tricyclic plasma concentration produced by these metabolic differences is substantial, and is often described as a **40-fold variation** between individuals. That is the pharmacological argument for measuring a level in this class rather than inferring one from the dose, and it is an argument about spread rather than about any particular target. The concentrations defining an adequate or a toxic level, and the framework for therapeutic drug monitoring, belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes; the enzymology itself belongs to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

IN INDIA

The pharmacogenetics of this class are not a footnote in Indian practice. About **20% of Asians** carry a polymorphism producing deficient CYP2C19 metabolism of the tertiary tricyclic compounds, and the reduced-clearance CYP2D6 allele above is carried by a large fraction of the same population. The consequence is a starting position, not a prohibition: begin at the low end, titrate against response and tolerability rather than towards a dose that suited someone else, and read a heavy anticholinergic, sedative or postural burden appearing at a modest dose as a probable metabolic phenotype rather than as intolerance to the class. One caution on the evidence: the source uses the broad category of Asian descent and reports no India-specific frequencies, so the direction is well grounded and the exact local proportion is not.

PITFALL

Choosing a tricyclic dose by analogy with the last patient who did well on one. The spread in plasma concentration between individuals on identical dosing is wide enough that a dose is a poor description of what a patient is actually receiving, and this is the class in which that gap is widest. A clinician who transfers a dose between patients, or who reads an unexpectedly heavy adverse-effect burden as the patient being difficult, is treating the dose as though it determined the exposure. The dose is an input; the response, the tolerability and where indicated the measured concentration are the output.

16.8 What the pharmacology decides, and what it does not

The class is dangerous for reasons that are entirely legible, and it survives for a reason that is equally legible. The source is explicit that what the tricyclics lost was not effectiveness; the charge against them is a set of actions bolted onto an effective molecule, not a weak molecule.

- Which monoamine an agent engages, from the methylation of its side chain, with metabolism blurring the distinction in the patient.
- Which adverse effects a patient will report, from the receptor actions accompanying the reuptake block rather than the block itself.
- Why the electrocardiogram is the relevant investigation, and why conduction is the parameter this class affects.
- Why a measured concentration adds information a prescribed dose cannot.

What the pharmacology does not settle is whether this patient should be given this drug. That judgment weighs the illness, what has already been tried and the risk the individual carries, and it belongs to the Mood disorders: pharmacotherapy notes. Pharmacology contributes one thing to it, and contributes it forcefully: in a class whose dispensed quantity is itself close to a lethal quantity, safety is not an afterthought to the prescribing decision but part of the same decision.

A tricyclic treats through reuptake inhibition and harms through everything else it does: muscarinic, histaminergic, alpha-adrenergic and voltage-sensitive sodium channel blockade, of which the sodium channel is the one that causes coma and seizures centrally and arrhythmia, cardiac arrest and death peripherally, and the lethal quantity is about a month of the prescription itself.

17 · Monoamine oxidase inhibitors: the interaction burden that defines the class

This is the only antidepressant class whose safety depends on what the patient eats. At the level of mechanism the monoamine oxidase inhibitors are the simplest antidepressants there are: they stop an enzyme from deaminating monoamines, and intrasynaptic amine concentrations rise. Everything difficult about them follows from three properties of that enzyme rather than from anything happening at a receptor. It exists in two isoenzyme forms with different substrate preferences, so which form a drug blocks decides whether it is an antidepressant at all. Most of the body's monoamine oxidase is not in the brain, and the peripheral pool works as a metabolic barrier between the gut and the systemic circulation. And when the block is covalent, function returns only when new protein has been made, which uncouples the duration of the drug's effect from the duration of the drug.

Read the class through the enzyme and its dietary rules, its waiting periods and its interaction burden all become derivable instead of memorised. What follows is that reading, from the isoenzyme split to the two structural escapes. Whether an MAOI suits a particular depressive episode belongs to the Mood disorders: pharmacotherapy notes. Recognising and treating a hypertensive crisis or a serotonin syndrome, and the cross-class interaction tables, belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

40 mg	10 mg	Up to 2 weeks	1 to 2 days
TYRAMINE DEFINING A HIGH-TYRAMINE MEAL, FASTED	FASTED TYRAMINE THAT CAN ALREADY PROVOKE A REACTION	FOR MONOAMINE OXIDASE ACTIVITY TO RETURN AFTER AN IRREVERSIBLE INHIBITOR	FOR THE SAME RECOVERY AFTER A REVERSIBLE INHIBITOR

17.1 Two isoenzymes, and what the split does and does not predict

Monoamine oxidase exists in two forms that differ in what they prefer to break down and in what inhibits them. **MAO-A** preferentially deaminates serotonin and noradrenaline and is the enzyme blocked by clorgyline and by moclobemide. **MAO-B** preferentially deaminates phenylethylamine and is the enzyme blocked by selegiline. Dopamine sits between them and is degraded about equally by both.

Two cautions attach to that tidy table and examiners like both. The first concerns noradrenaline: the Indian pharmacology standard groups it with serotonin on the MAO-A side, leaving dopamine as the only shared substrate, while the Comprehensive Textbook of Psychiatry makes dopamine and noradrenaline substrates for both forms and leaves serotonin as the only preferential MAO-A substrate. Both passages are cleanly written, so no answer should rest on the proposition that noradrenaline is an MAO-A substrate alone. The second concerns serotonin and

MAO-B. One standard pharmacology text describes MAO-B as effective against serotonin and dopamine; the psychopharmacology sources restrict it to trace amines such as phenylethylamine, with serotonin handled only at high concentrations. What survives every version is that serotonin is the secure preferential substrate for MAO-A and phenylethylamine the secure one for MAO-B.

The functional consequence is more secure than the substrate list, and it is where the teaching should be anchored. Blocking MAO-B on its own produces no antidepressant effect, because it leaves serotonin and noradrenaline metabolism untouched and, with MAO-A still working, allows almost no dopamine to accumulate. The complementary and load-bearing statement is that brain MAO-A must be substantially inhibited before an antidepressant effect occurs. That single fact organises the whole class, and it predicts in advance that any MAO-B inhibitor pushed to an antidepressant dose will have stopped being selective.

17.2 The enzyme outside the brain, and the barrier an MAOI deletes

Most teaching about antidepressants stays inside the synapse, and for this class that is the wrong place to look. **Tyramine** is produced continuously by bacterial action on dietary amino acids, both during the fermentation and ripening of food and within the gut itself. In an untreated person it never becomes a systemic problem, because monoamine oxidase in the gut wall and in the liver deaminates it before it can reach the circulation in any concentration that matters. That peripheral pool is doing nothing for mood: it is doing customs work, keeping ingested amines out of the arterial circulation.

An MAOI does not distinguish between the enzyme that raises intrasynaptic amines and the enzyme that guards the portal circulation, so the therapeutic action and the class hazard are one molecular event happening in two places. Elsewhere a drug's receptor profile predicts its adverse effects; here they are written in the enzyme's anatomical distribution instead.

RULE **The defining hazard of this class is peripheral, not central.** What an MAOI removes is the gut and hepatic barrier that normally destroys ingested amines before they reach the systemic circulation. Every characteristic danger of the class descends from that one deletion, which is why an MAOI is an interaction problem rather than a receptor problem, and why nothing in its binding profile forecasts its most dangerous events.

17.3 The pressor reaction, and the load it actually takes

The cheese reaction is a four-step chain, and setting the steps out separately shows where each of this class's rules attaches.

- Fermented foods, classically cheese, beer and red wine, are enriched in tyramine and to a lesser degree in other phenylethylamines.
- In an intact person the gut wall and the liver destroy that load on the way through, so almost none of it is absorbed systemically.
- With gastrointestinal and hepatic monoamine oxidase inhibited, a large ingested load is absorbed rapidly and reaches the systemic circulation in high concentration.

- Tyramine is not a receptor agonist but an **indirect sympathomimetic**: it is carried into sympathetic nerve terminals and displaces stored noradrenaline.

The last step is where the isoenzyme split becomes clinically decisive. Noradrenaline released by tyramine is ordinarily destroyed by MAO-A before it can do anything, and only when MAO-A is inhibited does the same release raise blood pressure. The reaction is therefore specific to MAO-A inhibition rather than to monoamine oxidase inhibition in general. When it does occur, the release of noradrenaline is massive and precipitous, and the hypertension can be severe enough to precipitate a myocardial infarction or a stroke. How that crisis is recognised and treated at the bedside belongs to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

The load required is smaller than most candidates assume. A **high-tyramine meal** is generally defined as 40 mg or more of tyramine in the fasted state, but in a patient taking an oral nonselective irreversible MAOI as little as 10 mg in the fasted state can provoke a hypertensive reaction. The fasted qualification is part of the number: the same quantity taken with a meal is better tolerated, so a threshold quoted without it is a different threshold.

17.4 Irreversibility, or waiting for protein rather than for clearance

The nonselective agents used in depression are **irreversible inhibitors**, forming a covalent complex that the enzyme cannot shed, and the pharmacokinetic consequence is the most counter-intuitive property of the class. The parent drug is excreted within 24 hours, yet monoamine oxidase activity takes up to 2 weeks to recover, because recovery depends on the synthesis and transport of new enzyme to monoaminergic nerve terminals rather than on elimination of the drug. A plasma level would therefore be useless: the patient's pharmacology is the state of the enzyme, and the enzyme has no plasma level.

Every waiting period in this class is a consequence of that arithmetic, including the one patients are least likely to observe. A low-tyramine diet has to be kept during treatment and for a further two weeks after the drug is stopped, which is exactly the time the enzyme needs to come back. The same source records that the reaction is itself rare, and rarity is why the tail of the restriction gets abandoned: nothing happened during treatment, so nothing is expected after it.

MNEMONIC

Enzyme out, enzyme in

Every interval in this class counts enzyme, never drug. The wait before a new agent is started is the time needed to make **new enzyme**; the diet that continues after stopping is that same interval seen from the patient's side. Ask when the enzyme comes back, not when the drug goes away.

TRAP **Assuming the dietary restriction ends with the last tablet.** It does not: the diet continues for a further two weeks, because what protects the patient is the enzyme, and the enzyme has not yet been replaced. Candidates err here because every other drug they know stops mattering once it is cleared, and clearance is irrelevant to this one.

17.5 The washout arithmetic, in both of its directions

Switching is where this class hurts people, and the arithmetic is not symmetrical. Coming off an MAOI, the governing quantity is enzyme resynthesis, and it is the same whatever comes next. A minimum two-week washout is required when switching from an MAOI to any other antidepressant class, and equally when switching from one MAOI to another. That second clause surprises people: when the block left behind is covalent, an MAOI-to-MAOI switch is not a cross-taper but a stop, a wait and a restart.

Where the next drug is an SSRI, the same interval is stated as **at least 14 days** after the irreversible inhibitor has been stopped, and the reason given is explicitly enzymatic: that period allows new monoamine oxidase to be synthesised. Going the other way, the governing quantity is not the enzyme but the elimination of the incoming drug and of its active metabolites, which is why the number changes with the agent. Fluoxetine is the case worth memorising: **at least five weeks** should pass between stopping it and beginning an MAOI, because Goodman and Gilman reason from a norfluoxetine half-life of one to two weeks when they set that interval. That figure is their stated premise for the washout and not a settled class number: the Comprehensive Textbook of Psychiatry gives a wider range for the same metabolite, and the comparative kinetics of the SSRIs among themselves are taken up where that class is compared.

The figure is not agreed across sources. The Maudsley Prescribing Guidelines give five to six weeks for the same switch. A shorter interval circulates in question banks drawn from a single geriatric answer key; it is the outlier, and it errs in the direction that hurts the patient. Teaching at least five weeks respects the floor that the two major sources share.

TRAP **Applying one interval in both directions.** The MAOI-first direction is governed by enzyme resynthesis and is the same for every successor drug. The MAOI-second direction is governed by elimination of the drug already on board, so it is agent-specific and, for the SSRI with the long-lived active metabolite, considerably longer. Candidates who memorise one number get one of the two directions wrong, and it is usually the dangerous one.

GOOD TO REMEMBER

Before prescribing any new psychotropic to a patient who has recently stopped an MAOI, establish the date of the last dose and count forward from it. A patient two days off an irreversible MAOI is, for prescribing purposes, still fully inhibited.

17.6 Reversibility: the escape from the diet that actually works

Moclobemide is a **reversible inhibitor of MAO-A**, and it may be prescribed without dietary precautions. Read carelessly this looks like a lesson about selectivity, and it is not: moclobemide is selective for exactly the isoenzyme whose inhibition creates the tyramine hazard, so if selectivity were protective it would be the most dangerous drug in the class rather than the safest. The protective property is reversibility. Because the inhibition is competitive, a rising tyramine load displaces the drug from the enzyme, so the substrate that would otherwise accumulate restores its own metabolism.

The washout figures prove the same point arithmetically. Full monoamine oxidase activity is restored within 1 to 2 days of stopping moclobemide, against the fortnight an irreversible agent needs, and the Maudsley switching table accordingly specifies only a twenty-four hour gap after moclobemide before an irreversible MAOI is begun. What separates those two figures is not the isoenzyme but the bond: a competitive block ends as the drug leaves, a covalent one only when new protein has been made.

How complete the protection is remains contested. The Indian pharmacology standard states flatly that potentiation of the pressor response to ingested amines is insignificant and that dietary restriction is not required. The Comprehensive Textbook of Psychiatry, writing on the same class in anxiety disorders, holds that restrictions can be less stringent but are still recommended at higher doses. Both agree on the mechanism; they differ on whether protection is complete. Treat the mechanism as settled and the practice as dose-dependent.

IN INDIA

Where an MAOI is genuinely under consideration and repeated, reliable dietary counselling cannot be assured, the reversible MAO-A inhibitor is the within-class choice that removes the problem rather than managing it, because reversible competitive block carries no dietary precaution at usual doses. That is a pharmacological argument, not a convenience one: an irreversible agent transfers a permanent responsibility to the patient's kitchen for the whole of treatment and beyond it.

Note also what the sources do not supply. The fermented-food examples in every major text are drawn from Western diets, and no source available here gives an Indian equivalent list or an Indian tyramine table. Counselling therefore has to be built from the mechanism, which is aged, fermented and bacterially ripened food of any cuisine, and not from a translated list of cheeses.

17.7 Selectivity and route: why selegiline is two different drugs

Selegiline demonstrates that isoenzyme selectivity is a property of a dose, not of a molecule. Oral selegiline at **5 to 10 mg/day** selectively inhibits MAO-B and is not effective in depression, while the antidepressant doses, not taught here, are precisely the ones at which selectivity is lost and a tyramine-restricted diet becomes necessary; the same Indian pharmacology text records the loss of enzyme selectivity at oral doses above twenty milligrams per day. Both follow from the rule that antidepressant efficacy requires substantial MAO-A inhibition: no drug is both MAO-B selective and an antidepressant.

Route of administration offers the second escape, and it works on the anatomy rather than on the chemistry. Transdermal selegiline avoids **first-pass metabolism** and is relatively sparing of the gastrointestinal MAO-A system, so brain enzyme is inhibited while the gut barrier is comparatively preserved. At the lower-strength patch, 6 mg/24 hours, tyramine restriction is not needed when the patch is used as an antidepressant, an indication owned by the Mood disorders: pharmacotherapy notes.

The higher-strength patches are contested. The Maudsley reports no hypertensive reactions even at the higher strengths, while recording that the manufacturer advises avoiding very high

tyramine foods at those strengths. Stahl's prescribing text argues that a far larger fasted tyramine load would be needed to react at the top strength, then states that studies are insufficient to be sure of safety, and elsewhere in the same book recommends food restriction at the higher transdermal doses. Where a source contradicts itself, the manufacturer caution stands.

AGENT OR ROUTE	NATURE OF THE BLOCK	ISOENZYME AFFECTED	TYRAMINE RESTRICTION	TIME TO FULL ENZYME ACTIVITY
Oral nonselective MAOI	Irreversible, covalent	Both isoenzymes	Required, and continued beyond the last dose	Up to two weeks after stopping
Moclobemide	Reversible, competitive	MAO-A	Not required at usual doses; still advised at higher doses	One to two days after stopping
Oral selegiline, low dose	Irreversible	MAO-B selective, and not antidepressant	Not applicable at this dose	As for any irreversible inhibitor
Oral selegiline, antidepressant dose	Irreversible	Selectivity lost	Required	As for any irreversible inhibitor
Transdermal selegiline, lower-strength patch	Irreversible	Selectivity lost, as at any antidepressant dose	Not required, because gut MAO-A is relatively spared	As for any irreversible inhibitor

17.8 Where the danger of this class actually sits

The dietary reaction is what the class is famous for, and it is not what most often harms people. Drug interactions are potentially more important clinically than the tyramine reactions, may be more common than the dietary interactions, and some of them can be dangerous or even lethal. That ordering should reorganise how the class is taught. The diet is visible, countable and patient-facing, so it is what gets counselled, examined and remembered; the interaction burden is invisible, sits in someone else's prescription, and is the one that kills.

Two structural features make it so. The enzyme is deleted rather than merely occupied, so the deletion persists through any drug started during those weeks; and monoamine oxidase is the disposal route for a wide range of amines, endogenous and pharmaceutical alike, so any agent that raises an amine has lost its brake. The specific pairs and the emergencies they produce belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

PITFALL

Spending the whole consultation on the food list. Clinicians counsel the diet thoroughly because it is concrete and teachable, then discharge the patient with no written instruction that nothing new, bought over the counter or prescribed by another department, is started without being checked against the MAOI. Dietary counselling is the easy half of the safety plan, and not the half with the higher body count.

Where gut MAO-A is inhibited, reversibility and not isoenzyme selectivity decides whether a patient needs a tyramine-restricted diet, and the two escapes both work by sparing that enzyme instead, one by selectivity and one by route; every waiting period in this class measures the resynthesis of enzyme, never the elimination of drug.

Mood stabilisers

18 · Lithium: the narrow-index drug and what monitoring buys

Almost every drug in this chapter is titrated against the patient; lithium is titrated against a number. That asymmetry is a pharmacological fact rather than institutional caution, and two properties of the molecule produce it. The distance between the concentration that treats and the concentration that harms is small enough that an ordinary prescribing decision, a week of vomiting or a new antihypertensive can carry a stable patient across it. And the prescription does not tell you where in that distance the patient sits, because the relation between milligrams swallowed and millimoles circulating is set by the kidney rather than by the pen. The serum concentration therefore stops being a laboratory refinement and becomes the only instrument reporting what the patient is actually exposed to. What follows develops lithium as a pharmacological object: what the ion does, how the body handles it, and what the number buys that clinical observation cannot. Lithium as a treatment for bipolar disorder belongs to the Mood disorders: pharmacotherapy notes, and the published target concentrations, with the monitoring framework built around them, belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

20 to 24 hours	28 to 48 hours	80 to 100 percent	10 to 14 hours
PERIPHERAL HALF-LIFE AT STEADY STATE	HALF-LIFE IN THE BRAIN	BIOAVAILABILITY, ANY PREPARATION	POST-DOSE SAMPLING WINDOW

18.1 Why this molecule needs a level at all

A drug earns routine blood-level monitoring when watching the patient is not a safe substitute for knowing the exposure, and lithium is the archetypal narrow-index drug because it fails that substitution twice over. Its **therapeutic index** is narrow: the gap between the concentration at which the ion works and the concentration at which it harms is too small to cross safely by guesswork. The derivation of the index, and the wider question of which psychotropics deserve a

level, are developed in the section of these notes on therapeutic index and dose-response. For lithium the index is not a comparative statistic but a constraint on every prescription.

The second failure is the more instructive. Even with a generous window you could not aim at it from the prescription, because the oral dose required to produce any given blood level varies widely between patients, and the reason it varies is renal clearance. Two adults of similar build, given the same milligrams of the same salt, sit at different concentrations, and neither history nor examination identifies which of them is high. For most psychotropics that between-patient variability is absorbed by a wide margin of safety and never has to be measured. Here the variability and the narrow margin arrive together, and the level is what closes the gap between them.

The pharmacological differences between the mood stabilisers are differences of exactly this kind: how much room the index leaves, what a level can report, and what constraint dominates prescribing. The accompanying figure sets them side by side; each agent is developed in its own section of these notes.

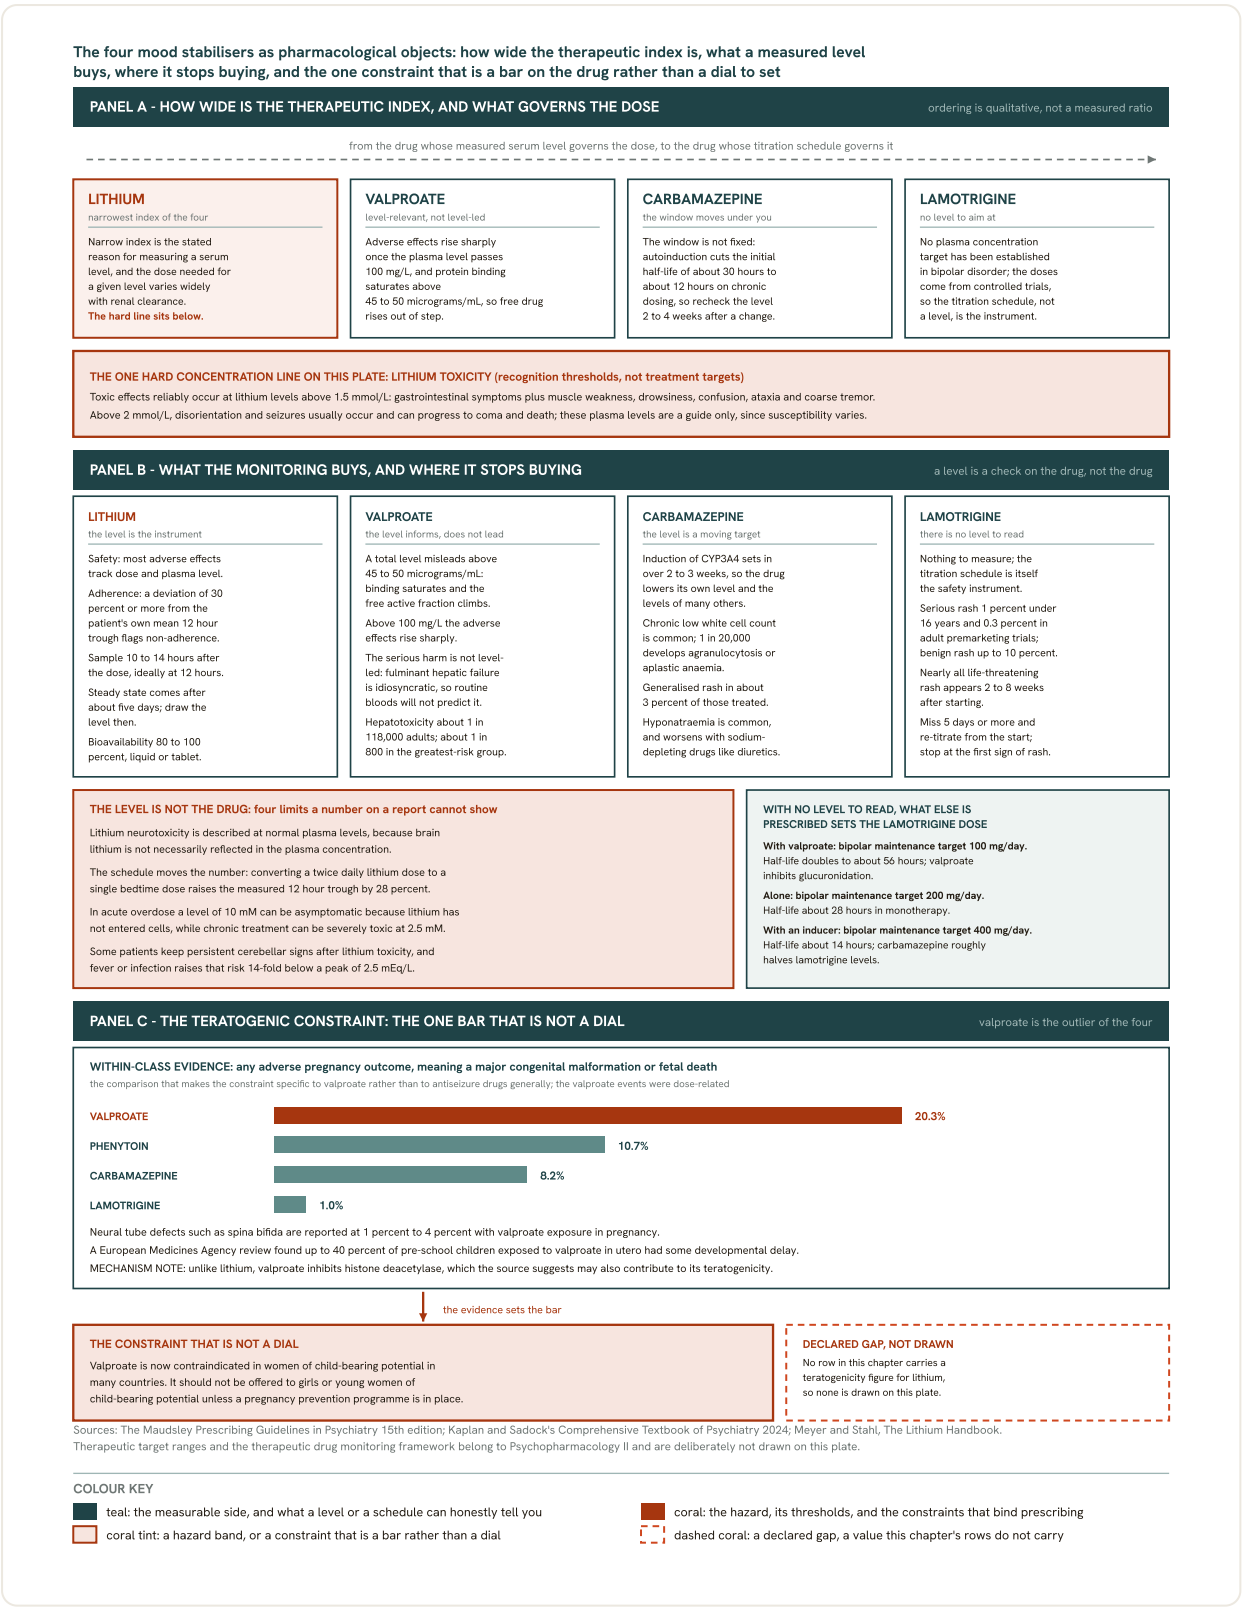

Fig 36.8 The four mood stabilisers compared as pharmacological objects. Lithium sits at the narrow end, where the serum level governs the dose and toxicity is reliable above 1.5 mmol/L; lamotrigine sits at the wide end, where no plasma target exists and the titration schedule is itself the safety instrument. A normal level does not exclude neurotoxicity, and the dosing schedule alone shifts the number by 28 percent. Valproate carries the one constraint that is a bar rather than a dial. (Maudsley 15th edition; Kaplan and Sadock, CTP 2024; Meyer and Stahl, The Lithium Handbook.)

RULE **A lithium dose is an instruction; a lithium level is a measurement, and only one of them describes the patient.** The dose tells you what you asked the kidney to handle; the level tells you what the kidney did with it. When they disagree, the level is right and the dose is a hypothesis.

18.2 A mechanism that has stayed plural

Lithium is the oldest drug in continuous psychiatric use and has the least settled mechanism of any agent in this chapter. That is the state of the science rather than a gap in the revision, and it is examinable as such. No single mechanism of action for lithium has been established, because the ion is implicated across so wide a range of biological processes that none has been isolated as the therapeutic route. The standard sources offer a shortlist of candidates instead, and the honest examination answer gives the shortlist and says plainly that it is one.

- **Glycogen synthase kinase 3**, an enzyme sitting at the junction of several intracellular signalling pathways.
- The cAMP response element-binding protein, which links signalling at the membrane to transcription in the nucleus.
- Pathways related to the sodium and potassium ATPase, the pump maintaining the ionic gradients lithium itself moves across.
- **Inositol monophosphatase**, the enzyme lithium inhibits inside the phosphatidyl inositol second-messenger system.

Two teaching points follow. All four candidates sit downstream of the receptor rather than at it. Every other class in this chapter is read forward from a receptor-binding profile: block this receptor, predict that adverse effect. Lithium cannot be read that way, which is why its adverse-effect signature has to be derived from distribution and renal handling rather than from a binding table. The biochemistry of the second-messenger cascades belongs to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

Second, an unsettled mechanism is not an unsettled drug: efficacy in mood stabilisation is established independently of any account of how it is produced. Candidates who present the mechanism as though one candidate had won lose marks, and so do those who present the uncertainty as though it undermined the indication.

18.3 Absorption, and what the preparations do not change

Almost nothing is lost between the tablet and the circulation. The **bioavailability** of lithium is 80 to 100 percent, and it does not depend on whether the preparation is liquid or tablet, or standard or sustained release. That is an unusually clean statement for an oral drug: switching preparations does not alter how much lithium reaches the blood over a dosing interval.

What the preparations differ in is the time taken to reach maximum concentration, which is the shape of the curve rather than the area beneath it. A sustained-release preparation flattens and delays the peak; it does not reduce the exposure. Two readings follow, pulling in opposite directions. A patient troubled by symptoms clustering shortly after the dose is a candidate for a formulation change, because the peak is what has moved. A patient whose total exposure is too

high is not helped by one, because the total exposure is what has not moved. Reaching for a sustained-release preparation to reduce exposure is a category error, and a common one.

18.4 Salt, element and ion: three numbers for one prescription

Lithium is prescribed as a salt, weighed as an element and acts as an ion, and those are three different quantities describing one tablet. The equivalence is the only route between a prescription written in one convention and a reference written in another. On the United States convention, 300 milligrams of lithium carbonate corresponds to 56 milligrams of **elemental lithium** and to 8 milliequivalents of lithium ion. United Kingdom formulations are described in millimoles instead: a 400 milligram lithium carbonate tablet carries 10.8 millimoles of lithium.

This is not pedantry. The amount of elemental lithium delivered depends on which salt the manufacturer used, so two products of nominally the same strength are not automatically interchangeable, and the Maudsley Prescribing Guidelines makes the consequence explicit: brand and strength both have to be specified, since the milligram figure alone does not identify the salt.

IN INDIA

Two decisions follow, and both are made while writing the prescription. First, name the brand and the strength rather than writing lithium and a milligram figure, because the elemental lithium delivered varies with the salt and a substitution at the counter is then a change of exposure. Second, read the unit the laboratory printed rather than the unit the textbook used: reports here follow the millimolar convention of the United Kingdom sources, while much of the reference literature is written in milliequivalents.

TRAP **Converting a lithium concentration between millimoles per litre and milliequivalents per litre.** The units genuinely differ: United Kingdom guidance reports lithium in millimoles per litre and the American texts in milliequivalents per litre or in millimolar, written mM, so a candidate who has read both has seen the same threshold written three ways and assumes a conversion factor exists. It does not. Millimolar is millimoles per litre under another name, and lithium is a monovalent ion, so all three are numerically identical and a figure quoted in one may be read directly as another. The error is not misreading the unit; it is multiplying by something to correct for it.

18.5 The kidney is the whole of the pharmacokinetics

Lithium is not metabolised in any way that matters to prescribing, so its disposition is a renal story from beginning to end. The cytochrome map that organises the interaction risk of almost every other agent here has no purchase on it, which is why an interaction with lithium is nearly always renal rather than enzymatic. The ion is freely filtered at the glomerulus, and about 70 percent of the reabsorption that follows takes place in the proximal tubule, where lithium is handled by **the sodium and hydrogen exchanger 3, or NHE3**, in direct competition with sodium.

That single sentence generates most of what a clinician has to anticipate. Because lithium and sodium compete for the same proximal transporter, the tubule cannot distinguish between them, so when sodium runs short and the proximal tubule turns avid, lithium is reabsorbed alongside it.

Sodium depletion, volume depletion and drugs acting on proximal handling therefore raise the level without a milligram being added to the prescription. A stable dose produces a stable concentration only in a stable patient, and the stability belongs to the patient rather than to the drug. The precipitants, and the recognition of an established toxic state, are developed in the section of these notes on lithium toxicity and the precipitants a prescriber controls; the renal, thyroid and parathyroid consequences of long-term exposure belong to the Endocrine and metabolic psychiatry notes.

18.6 Two half-lives, and the interval before a level means anything

Lithium has two half-lives, and confusing them produces confident errors in both directions. In adults at steady state the peripheral half-life is **20 to 24 hours**, while the half-life in the brain is longer, at 28 to 48 hours. A blood sample therefore reports the compartment that turns over quickly rather than the one in which the drug is acting, which is the kinetic basis for two observations that would otherwise look like folklore: the delay between an adequate concentration and an adequate effect, and the persistence of central effects after the serum concentration has come down.

Because the peripheral half-life is close to a full day, **steady state** is not reached quickly. Steady state arrives after five half-lives, which for lithium is about **5 days**, and a level is best drawn once it has been reached. A concentration measured earlier describes a system that has not finished responding, and it systematically understates where the patient is heading.

The central half-life also shapes how the drug should be given. Because the brain compartment turns over slowly, it is not exposed to the peaks and troughs the serum shows across a day. That long central half-life, taken together with the renal and adherence costs of dividing the dose, is the reason The Lithium Handbook gives for administering maintenance lithium as a single bedtime dose rather than in divided doses.

> **PITFALL**
>
> Drawing a level a day or two after a dose change and titrating on the result. The concentration is still climbing towards its new plateau, so the number is low, and the response to a low number is to increase the dose again. The patient then arrives at steady state on two increments instead of one, and the second answered an artefact of timing rather than a real deficit. After any change in dose, wait for the system to settle before asking it a question.

18.7 The trough is a convention, not a law of kinetics

Every published lithium concentration assumes a sampling time, and the assumption is so uniform it is rarely stated aloud. Blood for a lithium level is drawn **10 to 14 hours** after the last dose, ideally at 12 hours, in a patient taking a single bedtime dose of a prolonged-release preparation. That window is why a morning clinic and a bedtime dose fit together so conveniently, and why a sample taken at an arbitrary time cannot be compared with any published figure.

The interval has no special kinetic authority, and knowing why protects against over-interpreting it. **The twelve hour trough** became the standard because it was convenient rather than kinetically necessary, and the peripheral half-life is in fact closer to 24 hours. A drug whose half-life is around a day has not fallen far by that point, so the measured trough is not a nadir in any meaningful sense; it is a reproducible point on a slowly declining curve, and reproducibility is all a monitoring standard needs.

The corollary catches candidates and clinicians equally. If the number depends on the interval, changing the schedule changes the number without changing exposure. Converting a patient from twice daily lithium to a single bedtime dose raises the measured trough by **28 percent**, so the same total daily intake reads differently depending on how it is divided. Nothing about the exposure has moved; the sample now sits at a different point on the curve.

TRAP **Reading a rise in the trough after a schedule change as a rise in exposure.** The stem gives a patient whose total daily dose is unchanged, whose regimen has been consolidated into a single night-time dose, and whose next level is higher; the expected wrong answer is to reduce the dose. The increase is a property of when the blood was taken relative to the dose, and the correct move is to re-establish a baseline on the new schedule. The same logic condemns any comparison between two levels drawn at different intervals.

Set out together, the kinetic facts make a short table in which each line decides something.

KINETIC PROPERTY	GROUNDED FIGURE	WHAT IT DECIDES
Bioavailability	80 to 100 percent, independent of preparation	Switching preparation changes the peak, not the exposure
Proximal tubular reabsorption	About 70 percent of reabsorption	Sodium handling, not enzymes, drives interactions
Peripheral half-life	20 to 24 hours	How fast the serum compartment responds
Half-life in the brain	28 to 48 hours	Why effect lags and outlasts the serum figure
Time to steady state	Five half-lives, about 5 days	When a level after a dose change becomes meaningful
Sampling interval	10 to 14 hours post-dose, ideally 12	Whether the result can be compared with anything

18.8 What the level buys beyond safety

Monitoring is usually taught as a defensive activity, and that undersells it. Safety is the obvious purchase: most of the adverse effects of lithium are related to the dose and to the plasma level, so a patient's complaints can be interrogated as questions about exposure rather than treated as fixed properties of the drug. An adverse effect that tracks the concentration may respond to a change in the concentration, and that is a different conversation from the one ending in stopping the drug.

The second purchase is information no other instrument in psychiatry provides: an objective statement about whether the drug is being taken. Used naively that is crude, since between-patient variability means a single low value proves nothing. Used as a within-patient comparison it is informative, because a deviation of **30 percent or more** in a patient's twelve hour trough from that patient's own mean baseline level is a useful indicator of non-adherence. The comparator is the patient, not a population, which is why the first few levels on a stable regimen matter more than they appear to: they establish the baseline every later result is read against.

GOOD TO REMEMBER

Three questions make a lithium result interpretable, and asking them takes less time than the assay. When was the last dose taken, and how long before the sample was drawn. Is the patient on a single bedtime dose or a divided regimen, and has that changed since the last result. Has anything altered sodium or volume status since, including a new antihypertensive, a gastrointestinal illness or heavy sweating. A result without those three answers is a number, not a measurement.

MNEMONIC

Time since dose, Interval to steady state, Method of dosing, Estimate of adherence

Time since the last dose, which fixes where on the curve the sample sits. **I**nterval to steady state, which fixes whether the system had finished responding. **M**ethod of dosing, single or divided, which fixes what the trough means. **E**stimate of adherence, read against the patient's own baseline. All four are properties of how the sample was taken rather than of the drug.

None of this yields a target. The concentrations to aim for and the framework for acting on a result belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes. What this section establishes is why that apparatus exists: a narrow index leaving no room for approximation, a dose that does not predict a concentration, a kidney that alters the concentration without anybody altering the dose, and a sampling convention rigid enough that the number can be compared with something.

Lithium is measured because its therapeutic index is narrow and the dose needed to reach a given level varies with renal clearance; the number is interpretable only at steady state, at a fixed interval after the dose, on a known schedule, and read against that patient's own baseline.

19 · Lithium toxicity: recognition and the precipitants you control

Lithium toxicity is not an idiosyncratic reaction to the drug; it is the same drug at a higher concentration, and the concentration almost always rose because something happened to sodium. That sentence organises everything below, and it is why toxicity belongs in a systematic pharmacology chapter rather than only in an emergency one. Every other class in these notes generates its adverse-effect signature at a receptor, so the signature can be read forward from a

binding profile. Lithium offers no such profile. Its effects at high concentration are the effects of an ion that has accumulated, and accumulation is decided in the renal tubule by events that have nothing to do with psychiatry: a diarrhoeal illness, a new antihypertensive, a week of heavy sweating, an analgesic taken for a sprained ankle. The pharmacological task is therefore twofold: recognising the state and knowing what a plasma concentration can and cannot say about it, and then knowing in advance which prescriptions and which physiological changes carry a stable patient across the threshold. The management of the toxic patient, including escalation and the point at which extracorporeal removal is considered, belongs to the Mood disorders: pharmacotherapy notes. The renal and thyroid consequences of long-term exposure belong to the Endocrine and metabolic psychiatry notes, and the published target concentrations, together with the cross-class interaction matrix, belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

above 1.5 mmol/L TOXIC EFFECTS RELIABLY APPEAR	**above 2 mmol/L** SEIZURES; CAN PROGRESS TO COMA AND DEATH	**sevenfold** HOSPITALISATION RISK, ACE INHIBITOR IN THE ELDERLY	**14-fold** PERSISTENT CEREBELLAR SEQUELAE IF FEBRILE, PEAK BELOW 2.5 MEQ/L

19.1 The concentration at which the ion becomes a poison

Toxicity has a recognisable clinical shape and a concentration at which that shape becomes reliable. The Maudsley Prescribing Guidelines put it plainly: toxic effects occur dependably at concentrations **above 1.5 mmol/L**, and what appears there is a combination of gastrointestinal symptoms with central nervous system effects that includes muscle weakness, drowsiness, confusion, ataxia and **coarse tremor**. The two halves are not equally informative. The gastrointestinal component is non-specific and by itself as likely to be gastroenteritis. The neurological component makes the diagnosis, and its most useful feature is the character of the tremor: a fine tremor accompanies lithium at concentrations doing therapeutic work, whereas a coarse tremor with ataxia and slowed cognition beside it is a statement about exposure until proved otherwise.

Higher still, the picture ceases to be a matter of tolerability. **Above 2 mmol/L** disorientation increases and seizures usually occur, and the picture can progress to coma and to death. The escalation from a drowsy patient with a coarse tremor to a patient having seizures is therefore compressed into a very small distance on the concentration axis, which is what a narrow therapeutic index means in practice and why the account of lithium as the archetypal narrow-index drug sits alongside this one in these notes. The source qualifies both figures immediately, and the qualification is examinable in its own right: these plasma concentrations are a guide only, because individual susceptibility varies.

TRAP **Treating the toxicity threshold as a single number that holds for every patient.** The standard texts disagree about where toxicity begins, and the disagreement is genuine rather than an error in one of them. The Maudsley Prescribing Guidelines place the reliable appearance of toxic effects above the figure given here, which is written for the general adult patient. Geriatric Psychiatry Study Guide defines toxicity at a distinctly lower concentration and pairs that definition with a correspondingly lower target range in older patients. The two are not contradicting each other about lithium; they are describing different populations, and an older patient reaches harm at a concentration a younger one tolerates. The error is to blend them into one wide band, which yields a threshold too high for the elderly and too low for everybody else. Quote the figure with its population attached, and take the ranges themselves from the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes, which own them.

19.2 Why the number and the patient can disagree

Those thresholds are indispensable and they are also among the most over-read numbers in psychiatry, because a plasma concentration measures the wrong compartment. Lithium acts inside cells; the sample is drawn from the circulation. In equilibrium the plasma figure is an excellent proxy for exposure, and out of equilibrium it misleads in either direction. Both directions are standard examination material and they point to opposite errors.

The first is the more dangerous. Lithium neurotoxicity has been described at plasma concentrations that are entirely normal, because the concentration of lithium in the brain is not necessarily reflected in the plasma concentration. A patient who has become ataxic, dysarthric or confused on lithium therefore carries a clinical diagnosis that a reassuring laboratory report does not overturn. The report tells you what is in the serum; the signs tell you what is in the brain, and where they conflict the signs describe the compartment the drug is acting in. This is the strongest argument for examining a lithium patient rather than merely monitoring one.

The second direction runs the other way and is best seen by holding acute overdose against **chronic lithium toxicity**. In an acute ingestion a patient may arrive with a blood concentration as high as 10 mM, a concentration known to be highly toxic, and be completely asymptomatic, because the ion has not yet moved into the cells. That patient is not well; that patient is early, and the absence of signs is a timing artefact rather than reassurance. The mirror image is the patient who has taken lithium for years and drifted: severe toxicity, with diarrhoea and seizures, can appear at a much lower concentration of 2.5 mM, because the intracellular and blood concentrations are already in full equilibrium. The same laboratory value means opposite things depending on how long the ion has had to distribute, and the history rather than the assay supplies that information.

TRAP **Concluding that a normal lithium concentration excludes lithium toxicity.** The stem gives a patient on long-term lithium with new ataxia, coarse tremor or confusion, and a concentration inside the usual range, and invites the candidate to look elsewhere for a cause. The pharmacology forbids that inference, because the brain concentration is not obliged to follow the plasma concentration. The correct answer treats the clinical picture as primary and the assay as one input among several. The same logic condemns the opposite move, which is discharging an asymptomatic patient after an acute ingestion because they look well when the blood was taken.

19.3 One principle behind almost every precipitant

Learn the sodium principle and the precipitant list stops being a list to be memorised. Most of the risk factors for lithium toxicity are changes in the sodium level, or changes in the way the body handles sodium, and the categories the standard sources give fall out of that one statement:

- Low salt diets, whether prescribed for hypertension or adopted by the patient.
- Dehydration, and any state of **volume depletion**, from a diarrhoeal illness to prolonged sweating in hot weather.
- Drug interactions, which for this drug are nearly always renal events rather than enzymatic ones.
- Some uncommon physical illnesses that disturb sodium handling.

The mechanism that unifies them sits in the proximal tubule, and every precipitant below is an application of it. Lithium and sodium share the same proximal tubular exchanger, so any medication or physiological state that lowers the serum sodium causes lithium to be preferentially reabsorbed, and the concentration climbs with no change to the prescription. The tubule cannot tell the two ions apart, so when sodium becomes scarce the kidney does what it is built to do, retaining sodium avidly and retaining lithium in the same movement. That is why volume depletion from any cause deserves the same attention as a new prescription.

RULE **Almost every precipitant of lithium toxicity is a sodium event, so ask what has happened to sodium before asking what has happened to the dose.** The question sorts the whole field: a new antihypertensive, a diuretic, an anti-inflammatory, a gastrointestinal illness, a fast, a heat wave and a drug that causes hyponatraemia all arrive at the same tubular endpoint by different routes. It also tells you where to look when a concentration has risen and the prescription has not changed, which is the commonest clinical version of this problem. The answer is rarely in the drug chart column headed lithium.

19.4 Volume, and why a thiazide is not a loop diuretic

Diuretics give the cleanest demonstration, because two classes that both make a patient pass urine carry markedly different lithium risk, and the difference is entirely where in the nephron each acts. The thiazide is the dangerous one. Lithium concentrations usually rise **within 10 days** of a thiazide diuretic being started, and the magnitude of the rise is unpredictable, varying from an increase of 25 percent to 400 percent. A range that wide is not a failure of the evidence; it is the claim. No dose adjustment can be planned in advance, and the interaction has to be measured in the individual patient rather than anticipated from a table.

The loop diuretic behaves differently, and the reason is anatomical. Only 10 percent of lithium is reabsorbed in the ascending thick loop of Henle, by way of **the sodium, potassium and chloride cotransporter NKCC2**, and blocking that transporter with a loop diuretic such as furosemide presents a risk of lithium toxicity only in elderly individuals. The loop is a minor site of lithium handling, so a drug that blocks it removes a small part of the reabsorptive machinery, whereas a thiazide acts distally and provokes the compensatory proximal sodium avidity that drags lithium with it. That is the site-of-action reason behind the ranking every text gives, and it

is more useful than the ranking itself: thiazides reduce the renal clearance of lithium more than loop diuretics do.

Safer is not safe, and the ranking is about the drug rather than the patient's fluid state. Anyone rendered genuinely hypovolaemic by a diuretic is at risk whatever its class, because the precipitant is then the volume rather than the transporter.

19.5 Three prescriptions that raise a concentration without changing the dose

Beyond the diuretics, three groups account for most of the interaction risk a lithium patient carries, and none of them is a psychotropic decision. Angiotensin-converting enzyme inhibitors come first, because they combine a large effect with a slow one. ACE inhibitors produce an unpredictable rise in the lithium concentration, of up to fourfold, that develops over several weeks, and in elderly patients they raise the risk of hospitalisation for lithium toxicity sevenfold. Angiotensin II receptor antagonists may be associated with a similar risk, so switching within the renin-angiotensin blockers is no solution. The slow onset is what catches clinicians: a concentration checked promptly after the new prescription can be entirely normal and mean nothing at all.

Non-steroidal anti-inflammatory drugs come second, and are pharmacologically the most instructive of the three. NSAIDs inhibit the synthesis of **renal prostaglandins**, thereby reducing renal blood flow and possibly increasing the renal reabsorption of sodium and therefore of lithium. The effect is again unpredictable in size and in timing: reported rises range from around 10 percent to over 400 percent, with an onset that varies from a few days to several months. What makes this class clinically dominant is how it is taken. These are drugs used on a when-necessary basis and bought without a prescription, so they enter the patient's regimen without entering the drug chart, and a review that asks only about prescribed medicines will not find them.

Third are the drugs that act on lithium at one remove, by moving the serum sodium itself. Carbamazepine can cause **hyponatraemia**, which in turn leads to lithium retention and toxicity, and rare reports of central nervous system toxicity implicate the selective serotonin reuptake inhibitors for the same reason. This is the one route by which a psychotropic co-prescription becomes a lithium risk, and the mechanism is not enzymatic at all. It is recorded here as a sodium-handling event; the systematic cross-class interaction matrix belongs to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

AGENT OR CLASS	RENAL MECHANISM	EFFECT ON THE LITHIUM CONCENTRATION	TIME COURSE
ACE inhibitors, and probably angiotensin II receptor antagonists	Reduced renal perfusion with sodium retention	Unpredictable, up to fourfold	Develops over several weeks
Thiazide diuretics	Distal sodium loss with compensatory proximal reabsorption	Unpredictable, an increase of 25 percent to 400 percent	Usually within 10 days
Loop diuretics	Block NKCC2 in the ascending thick loop of Henle, where only 10 percent of lithium is reabsorbed	Smaller than a thiazide	Risk described in elderly individuals
NSAIDs	Inhibit renal prostaglandin synthesis, reducing renal blood flow	Around 10 percent to over 400 percent	A few days to several months
Carbamazepine and the SSRIs	Hyponatraemia, leading to lithium retention	Indirect, by way of the serum sodium	Variable

PITFALL

Rechecking a concentration too early after a new interacting drug and reading the normal result as an all-clear. The onset of these interactions is measured in weeks for a renin-angiotensin blocker and can run to several months for an anti-inflammatory, so a level drawn a few days after the new prescription samples a system that has not finished moving. The result is not wrong, but it answers a question nobody asked. Check when the interaction has had time to express itself, and check again the moment the patient reports neurological features rather than waiting for the next scheduled test.

19.6 What an episode of toxicity can leave behind

Toxicity is usually described as though recovery were the default and the only question were how fast the concentration can be brought down. That is optimistic. Some patients retain persistent cerebellar signs after an episode of lithium toxicity, and the picture has a name of its own, **the syndrome of irreversible lithium-effectuated neurotoxicity**. The residual signs are cerebellar rather than diffusely cognitive, so a patient who remains dysmetric or ataxic once the concentration has normalised is not simply slow to recover.

The risk factor most easily missed is not the peak concentration but the intercurrent illness beside it. Fever or infection increases the risk of persistent cerebellar sequelae **14-fold** in patients whose peak serum concentration stayed below 2.5 milliequivalents per litre. The population that figure is drawn from is the point of it: these are not the patients with the most extreme concentrations, and the multiplier applies precisely to the group whose peak concentration would not by itself have predicted lasting damage. The source states the concentration in milliequivalents per litre, which for the monovalent lithium ion is numerically

identical to millimoles per litre, as the account of lithium as a narrow-index drug elsewhere in these notes sets out. The febrile, intercurrently unwell patient who becomes toxic is therefore the one to worry about most, and the acute management of that patient belongs to the Mood disorders: pharmacotherapy notes.

19.7 The precipitants a prescriber actually controls

The sodium principle converts a long and unmemorable list into a short set of questions asked at every review. Some precipitants are physiological and can only be anticipated; others arrive on a prescription pad and are genuinely controllable. The controllable ones cluster tightly and are worth holding as a group.

MNEMONIC

ACE inhibitors, NSAIDs, Thiazides, Sodium loss

ACE inhibitors, and the angiotensin II receptor antagonists beside them. **N**SAIDs, including the ones the patient bought rather than the ones you wrote. **T**hiazides, which are the diuretic class that matters most here. **S**odium loss of any kind, whether from a low salt diet, from a diarrhoeal illness, or from a drug such as carbamazepine that lowers the serum sodium directly. ANTS suits a group of small events that arrive quietly and add up, and every letter is a change in sodium handling rather than in the lithium prescription.

GOOD TO REMEMBER

Three questions turn the pharmacology into a review that takes under a minute. First, has another prescriber added anything, with particular attention to blood pressure treatment, diuretics and analgesia, and has the patient bought anything over the counter. Second, has anything happened to fluid or salt intake: a gastrointestinal illness, fasting, a new dietary restriction, a spell of heavy sweating. Third, is the patient febrile or intercurrently unwell, which changes both the risk of toxicity and the risk that toxicity leaves something behind. A patient who answers no to all three and whose concentration has risen has an incomplete history, not a mysterious kidney.

Recognising toxicity and treating it are different tasks, and this section owns only the first: the kinetics by which an ordinary prescribing decision, made for a good reason by somebody who was not thinking about lithium, moves a therapeutic concentration into a toxic one.

Lithium toxicity is a renal event before it is a neurological one: almost every precipitant is a change in sodium or in sodium handling, the plasma concentration is a guide rather than a verdict because the brain compartment need not follow it, and the prescriptions that matter most are the ones written by somebody else.

20 · Valproate: efficacy against the teratogenic prohibition

Valproate is the one agent in this class whose pharmacology has to be learned alongside a prohibition. Nothing in its chemistry, its absorption or its tolerability at ordinary concentrations explains why an entire population of patients is now excluded from receiving it. That exclusion

rests on one body of evidence about what the molecule does to a developing fetus, and it is why its formulary entry reads as a restriction before an indication.

What follows treats valproate as a member of a drug class rather than as a treatment for anybody: mechanism, kinetics, the relation between concentration and harm, its two kinds of adverse reaction, the comparative teratogenicity that separates it from the other antiseizure mood stabilisers, and its behaviour as a metabolic inhibitor. Its efficacy in mania and in maintenance belongs to the Mood disorders: pharmacotherapy notes. Its use as an antiepileptic, and the teratogenicity question as it arises in a woman treated for epilepsy, belong to the Epilepsy and psychiatry notes. Target concentrations, the cross-class interaction matrix and the framework for prescribing in pregnancy belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

3 to 4 hours	10 to 16 hours	20.3 percent	up to 40 percent
TIME TO PEAK PLASMA CONCENTRATION	PLASMA HALF-LIFE IN ADULTS	ANY ADVERSE PREGNANCY OUTCOME	DEVELOPMENTAL DELAY AFTER EXPOSURE IN UTERO

20.1 A mechanism list, not a mechanism

Chemically this is one of the least elaborate molecules in psychiatry. Valproate is a **simple branched-chain fatty acid**, and that ordinariness is the first clue to why its pharmacology has never resolved into one account. A short fatty acid has no obvious receptor to occupy, and the sources offer instead a set of actions at several levels of the cell, none isolated as the therapeutic route.

- Inhibition of the catabolism of gamma-aminobutyric acid, raising inhibitory tone by slowing the transmitter's breakdown rather than by acting at its receptor.
- Activation of the extracellular signal-regulated kinase pathway, a signalling route linked to synaptic plasticity.
- Promotion of brain-derived neurotrophic factor expression, placing valproate among the agents credited with neurotrophic effects.
- Reduction of protein kinase C levels, a second signalling action.

Two further routes are proposed rather than established, and the distinction matters. Depletion of inositol, and indirect effects on non-GABA pathways through the inhibition of **voltage-gated sodium channels**, are put forward as further proposed mechanisms, without either being shown to be the therapeutic route. The inositol action is the interesting one for this class, because valproate shares inositol depletion with lithium: two structurally unrelated drugs converge on one second-messenger substrate. The biochemistry of those cascades belongs to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

Most of this chapter reads a drug forward from its receptor-binding profile: block this receptor, predict that adverse effect. Valproate has no binding table to read, so its adverse-effect signature

has to be derived from its plasma binding, from what its concentration does as it rises, and from a class of reactions unrelated to concentration.

RULE **Valproate produces harms of three kinds, and only two of them answer to anything measured in the patient.** Some track the concentration and improve when it falls. Some appear unpredictably in a small number of exposed patients and are anticipated by no level and no blood test. And one falls on a person who is not the patient and is fixed at the moment of exposure; the reported events were dose-related, so dose changes its size, but no dose is named at which it stops and nothing in the patient reports it. Every prescribing constraint on this drug follows from holding those three apart.

20.2 Histone deacetylase: the action lithium does not have

One action separates valproate cleanly from the other archetypal mood stabiliser, and it is the most examinable line in the mechanism. Unlike lithium, valproate inhibits **histone deacetylase**, an effect that may account for many of its actions on gene transcription. Deacetylation tightens chromatin and silences transcription; inhibiting the enzyme leaves chromatin more open and expression more permissive. A drug acting at this level is not modulating a synapse, it is modulating which genes are read.

The same source goes one step further, and the step must be reported with its uncertainty intact. It is also possible that histone deacetylase inhibition contributes to valproate's teratogenicity. That is a hypothesis offered by the source, not an established causal pathway. Its interest is structural: it is the only proposal in the mechanism list connecting how the drug works to its worst harm. A candidate who states it as fact has overstated the source; one who omits it has missed the topic's only mechanistic bridge.

20.3 Absorption, half-life and what an enzyme inducer does to them

The kinetics are straightforward, and two practical constraints follow from them. All valproate formulations are rapidly absorbed, with peak plasma concentrations reached at 3 to 4 hours, and the plasma half-life in adults is 10 to 16 hours. Rapid absorption means the formulation changes the shape of the curve rather than the total exposure. A half-life of that length means the concentration is not flat across a dosing interval, which is why several of this drug's adverse effects are described as peak-related rather than simply dose-related.

The half-life is not a fixed property of the patient. It is shortened when a **cytochrome P450 inducing drug** is co-prescribed, so a stable patient can lose exposure with no change to the prescription. Within this class that matters immediately, since the other antiseizure mood stabiliser in common use runs in the opposite metabolic direction; enzyme induction is developed in the section of these notes on carbamazepine and oxcarbazepine.

20.4 Saturable protein binding, and why the reported total misleads

This is the subtlety in valproate pharmacokinetics, and it is skipped in revision. Valproate is highly bound to albumin, and its binding sites remain unsaturated below serum concentrations of **45 to 50 micrograms per millilitre** but become saturated above that, so the pharmacologically active **unbound fraction** rises disproportionately. Above that region the albumin has no further

capacity and additional drug has nowhere to go but into the free compartment, while the laboratory continues to report bound and unbound together.

GOOD TO REMEMBER

When a valproate concentration is reported, ask what it is a total of before deciding what it means. Below the saturation region the figure tracks the active drug closely enough to serve as a proxy; above it the reported total understates how much free drug an increment has added. Increments therefore matter less than where the patient already sits, and adverse effects disproportionate to a small dose change are not necessarily exaggerated.

20.5 Two thresholds, two conventions, and neither of them a target

Two concentrations are quoted about valproate and they are easy to conflate, partly because they arrive in different units. The first is the binding threshold above; the second is a harm threshold. Many adverse effects of valproate are dose-related and often peak-level related, and they increase sharply in frequency and severity once the plasma concentration exceeds **100 milligrams per litre**. That is a statement about where tolerability deteriorates, not about where efficacy begins.

The unit problem is genuine and better stated than resolved by preference: one major source reports valproate in milligrams per litre, the other in micrograms per millilitre. Milligrams per litre is numerically identical to micrograms per millilitre, so the two thresholds read directly against each other. Side by side they say something useful: binding saturates well below the concentration at which adverse effects escalate, so a patient climbing towards the harm threshold has already spent some distance in the region where the free fraction outruns the total.

TRAP **Reading either valproate threshold as a therapeutic target range.** The stem quotes a concentration and asks what should be done next. Both figures are boundaries of quite different kinds: one marks where albumin binding saturates and the free fraction begins to rise disproportionately, the other where adverse effects escalate. Neither names a concentration to aim for. The second half of the trap is the unit convention, which differs between the sources and tempts candidates into a conversion that does not exist. Target concentrations belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

20.6 Dose-related harm and idiosyncratic harm behave differently

The most useful organising move here is to sort valproate's adverse effects by whether the concentration has anything to do with them. The dose-related group is the larger. Valproate causes dose-related tremor in up to a quarter of patients, usually an **intention or postural tremor**, a frequency high enough that it should be asked about rather than waited for. A very small proportion develop parkinsonism with cognitive decline, which reverses on discontinuation, and the reversibility is what to carry away: a parkinsonian picture in a patient on this drug is a question about the drug before it is one about a degenerative diagnosis.

The idiosyncratic group is far smaller and behaves differently. Valproate may very rarely cause **fulminant hepatic failure**, with young children receiving multiple anticonvulsants at greatest risk. The estimated risk of hepatotoxicity in adults is about 1 in 118,000, rising to about 1 in 800 in the profile the source names as greatest risk, which is polypharmacy, age under two years and

intellectual disability. Those figures come from a black-box warning table adapted from an older reference work and should be quoted as dated estimates. What survives the dating is the shape: a rare event in adults, concentrated by orders of magnitude into a young, multiply medicated, developmentally impaired group. Hyperammonaemia and pancreatitis belong to the same idiosyncratic family, rare and unpredictable, with no incidence figure quoted for either.

HARM CLASS	EXAMPLES GROUNDED HERE	RELATION TO CONCENTRATION	WHAT MONITORING CAN DO
Dose-related	Tremor, and the reversible parkinsonian picture	Dose-related and often peak-level related	A level explains the complaint and predicts the response to a reduction
Idiosyncratic	Fulminant hepatic failure, hyperammonaemia, pancreatitis	Not related to concentration	No routine test reliably anticipates the event
Developmental	Neural tube defects, developmental delay after exposure in utero	Dose-related in the reported series for malformation and fetal death; no dose relation is stated for developmental delay	Nothing measured in the patient reports the risk to the fetus

MNEMONIC

Dose-related, Idiosyncratic, Developmental

Dose-related harm answers to the concentration and can be negotiated with by lowering it. **I**diosyncratic harm does not, and is not anticipated by routine testing, so the defence is recognition rather than surveillance. **D**evelopmental harm is not a harm to the person swallowing the tablet at all, which is why no level and no liver test in the patient reports it; the reported events were dose-related, so a lower dose changes how large it is, and only non-exposure removes it.

PITFALL

Treating a schedule of liver function tests as protection against the hepatic risk. The serious hepatic event is idiosyncratic rather than dose-related, which is what makes routine monitoring an unreliable predictor of it: a normal result last month says nothing about this month, and the reassurance it generates is the problem. The same applies to hyperammonaemia, where a level checked because a protocol requires it is a different thing from one checked because a patient has become drowsy, confused or vomiting. Test in response to a change in the patient rather than to the calendar, and tell the family what change should bring them back.

20.7 The teratogenic prohibition and the comparison that grounds it

Everything above concerns the person taking the drug. The restriction that now defines valproate concerns someone else, and it rests on two harms. The first is structural: exposure in pregnancy carries an increased risk of **neural tube defects** such as spina bifida, reported at 1 percent to 4 percent. The second is neurodevelopmental. A review of studies by the European Medicines

Agency found that up to 40 percent of pre-school children exposed to valproate in utero experienced some form of developmental delay, with outcomes poorer than after carbamazepine or lamotrigine exposure. A malformation risk in the low percentages and a developmental risk reported at up to two in five of the pre-school children exposed are not the same order of statement, and it was the second that moved the regulatory position.

What makes the restriction specific to this molecule rather than to antiseizure drugs generally is the within-class comparison. The rate of any adverse outcome, meaning a major congenital malformation or fetal death, was 20.3 percent with valproate against 10.7 percent with phenytoin, 8.2 percent with carbamazepine and 1.0 percent with lamotrigine, and the valproate events were dose-related.

ANTISEIZURE AGENT	ANY ADVERSE PREGNANCY OUTCOME	POSITION WITHIN THE COMPARISON
Valproate	**20.3 percent**	Highest rate of the four
Phenytoin	10.7 percent	About half the valproate rate
Carbamazepine	8.2 percent	Lower again, an alternative in class
Lamotrigine	1.0 percent	Lowest of the four

Read that table as the argument rather than as a list. If every antiseizure drug carried a comparable risk, the response would be a class-wide caution. Because one agent sits several times above its nearest comparator and an order of magnitude above the lowest, the response is specific to that agent, and the within-class alternative is what makes a restriction workable. Valproate is now contraindicated in women of child-bearing potential in many countries, and the attached framework is explicit about the condition of any use: valproate should not be offered to girls or young women of child-bearing potential unless a **pregnancy prevention programme** is in place. Both are the United Kingdom and European regulatory position as reported by the standard prescribing guidance, and the clinical detail of operating such a programme belongs to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

IN INDIA

No position issued by an Indian regulator is asserted here, and a prescriber should establish the current local labelling rather than assume it matches the position above. What travels unchanged is the drug-level evidence, because a malformation rate, a developmental outcome and a within-class ranking are properties of the molecule and not of the jurisdiction. Three decisions follow at the point of writing the prescription, none depending on which regulator has ruled. Establish first that the indication could not be met by an agent whose comparative pregnancy outcome is lower. Second, if valproate is still the answer for a girl or a woman who could become pregnant, treat effective contraception and a recorded discussion of the teratogenic and developmental evidence as conditions of the prescription rather than good practice around it. Third, document that reasoning.

TRAP **Using the epilepsy-versus-bipolar uncertainty to soften the prohibition.** Almost all of this evidence was collected in women treated for epilepsy, and Kaplan and Sadock's Comprehensive Textbook of Psychiatry reports a cohort in which, among the small number of exposed women without an epilepsy diagnosis, the risks were not increased and were equal to those of unexposed controls; the same source notes that the risks appear more prominent in women with epilepsy than in those with bipolar disorder. Candidates read that as permission to prescribe. It is not. The Maudsley Prescribing Guidelines treats valproate as an established human teratogen and records the contraindication without reference to the indication, while warning that data drawn from epilepsy studies may not be precisely relevant to use in mental illness, which argues for caution in both directions rather than reassurance in one. Name the limitation as a limitation of how the evidence was gathered, and keep the rule.

20.8 Valproate as an inhibitor: the glucuronidation problem

The last piece of class pharmacology is what makes valproate a hazard to other drugs. Valproate raises the plasma levels of several drugs by inhibiting **glucuronidation**, lamotrigine among them. Glucuronidation is a conjugation route rather than an oxidative one, which is why this interaction is missed by clinicians who screen a new prescription only against the cytochrome map: a drug of no interest to the oxidative enzymes can still have its clearance impaired by a conjugation inhibitor.

The consequence is why the two drugs are so often examined together. A drug cleared largely by conjugation accumulates when valproate is added, and the response is a change to that drug's titration rather than to valproate's. The arithmetic for lamotrigine is developed in the section of these notes on lamotrigine and the titration that prevents the rash. The wider cross-class interaction matrix belongs to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

Set the profile out together and the shape is clear: mechanistically unresolved, kinetically well behaved apart from a saturable binding step, tolerable in most patients, dangerous unpredictably in a very small number, and hazardous to a fetus in a way no level and no liver test will detect. The contested question is not its efficacy but who may be exposed to it at all, and the teratogenic prohibition is the answer regulators reached once the within-class comparison showed the harm belonged to this agent rather than to its class.

Valproate's dose-related harms can be negotiated with by lowering the concentration and its idiosyncratic harms are anticipated by no routine test; its teratogenic and developmental harms fall on someone who is not the patient and are fixed at the moment of exposure, which is why the teratogenic prohibition restricts who may receive the drug rather than cautioning how it is given.

21 · Carbamazepine and oxcarbazepine: the enzyme-induction problem

Almost every interaction in this chapter is a transaction between two molecules; the defining interaction of carbamazepine is with itself. The drug increases the activity of the enzymes that clear it, so the concentration produced by a fixed prescription falls during the first weeks of

treatment without anybody altering the dose, and a patient adequately treated early in the course may be underexposed a few weeks later. Every other property that makes this pair examinable radiates from that single fact. The same enzyme induction that shortens the drug's own residence in the body shortens the residence of a great many co-prescribed compounds, so a stable regimen destabilises around it. And because induction is a physiological change rather than a chemical event, it takes time to develop and takes time to reverse, which means stopping the drug is as much a pharmacokinetic manoeuvre as starting it.

Oxcarbazepine belongs in the same section because it is the structural answer to that problem, and the answer is partial rather than complete. What follows develops the two agents as pharmacological objects: the channel actions they share, the potent induction only one of them has, the time course of that induction in both directions, and the adverse-effect signature that owes nothing to enzymes. Carbamazepine as a treatment choice in bipolar disorder belongs to the Mood disorders: pharmacotherapy notes, and its use as an antiepileptic to the Epilepsy and psychiatry notes. The enzymology of the cytochrome system itself is developed in the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes, while the cross-class interaction matrix, the target concentration ranges and the monitoring framework built on them belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

30 hours INITIAL PLASMA HALF-LIFE	**12 hours** HALF-LIFE ON CHRONIC DOSING	**2 to 3 weeks** TIME FOR AUTOINDUCTION TO DEVELOP	**One in 20,000** AGRANULOCYTOSIS OR APLASTIC ANAEMIA

21.1 Two molecules, one channel, and what the derivative adds

Carbamazepine acts principally on the membrane rather than at a receptor. It blocks **voltage-dependent sodium channels** and so inhibits repetitive neuronal firing, reduces glutamate release and decreases the turnover of dopamine and noradrenaline. That ordering matters for teaching: the monoamine effects are consequences of a channel action, not a separate binding profile, which is why this drug cannot be read forward from a receptor table the way an antipsychotic can. Its adverse-effect signature has to be derived from what the molecule does to excitable membranes and to hepatic enzymes, and not from occupancy at H1, M1 or D2.

Oxcarbazepine is a structural derivative of carbamazepine and keeps the channel action, but it is not simply a cleaner version of the same molecule. It works through three actions rather than one.

- Blockade of voltage-dependent sodium channels, which it shares with the parent compound.
- An increase in potassium conductance, which stabilises the membrane by a route carbamazepine does not use.
- Modulation of high-voltage activated calcium channels, altering the calcium entry that follows depolarisation.

The second structural point is about identity rather than mechanism. Oxcarbazepine acts as a pro-drug for **licarbazepine**, and eslicarbazepine, where it is available, is likewise a pro-drug for

licarbazepine. Three named agents therefore converge on one active molecule, which is worth holding because it explains why their pharmacodynamic profiles resemble one another so closely, and why a difference between them is far more likely to be kinetic than pharmacodynamic.

21.2 Autoinduction: the half-life that halves itself

Carbamazepine is a potent inducer of the cytochrome P450 oxidase enzymes, and particularly of CYP3A4. It is also a substrate of the system it induces, and those two facts together produce **autoinduction**: by increasing the amount of enzyme available, the drug increases its own clearance and reduces its own concentration. The effect is not subtle and it is not a curiosity. An initial plasma half-life of **around 30 hours** is reduced to **around 12 hours** on chronic dosing, so the elimination half-life of a patient stabilised on carbamazepine is substantially shorter than the half-life the same patient had on the first dose.

The time course is the part candidates omit, and it is the part that decides clinical behaviour. Enzyme induction requires new protein to be synthesised, so it is not immediate: autoinduction develops after 2 to 3 weeks of treatment. The consequence is a moving target. Steady state is a concept that assumes fixed clearance, and here clearance is not fixed during precisely the interval in which the clinician is trying to establish a maintenance dose. A patient can pass through an apparently satisfactory plateau early and then drift downwards as the enzyme population expands beneath them.

RULE **With carbamazepine the dose is fixed and the clearance is not, so a falling concentration on an unchanged prescription is the expected pharmacology rather than a deviation from it.** The question to ask of an unexpected carbamazepine result is therefore not what the patient did, but where in the induction time course the sample was taken, and how far the enzyme population had moved by then.

21.3 What autoinduction does to measurement

Because clearance changes during the first weeks, a concentration is only interpretable when it is anchored to a point in that time course. Plasma levels should therefore be checked **2 to 4 weeks** after starting carbamazepine or after an increase in dose, to confirm that the desired level is still being obtained. The interval is chosen to sit beyond the period in which induction develops, so that the number describes the induced state the patient will live in rather than the transient pre-induction state.

Separate two readings of that rule. A level drawn before induction is complete is not a wrong number, but it is a number about a system that has not finished changing, and it systematically overstates where the patient will settle. A level drawn after induction is complete is the one against which a maintenance dose can honestly be set. The distinction also explains why a dose increase resets the clock: the increase re-enters a period in which the enzyme population and the substrate load are adjusting to each other, and a sample taken immediately afterwards reports neither the old state nor the new one. The published concentration ranges themselves, and the wider framework of therapeutic drug monitoring, belong to the Psychopharmacology II:

emergencies, resistance, monitoring and interactions notes; what this section establishes is why a carbamazepine level has to be timed at all.

PITFALL

Reading a fall in the carbamazepine concentration during the first month as evidence of non-adherence. The clinician sees an adequate early result, a lower result a few weeks later, an unchanged prescription and an unchanged patient, and reaches for the explanation that fits almost every other drug. Here the drug itself has changed the denominator. The corrective is procedural rather than clever: record when treatment started, record when the dose last changed, and read every concentration against those two dates before reading it against the patient.

21.4 Induction as an outward-facing problem

The enzyme that carbamazepine induces does not work only on carbamazepine. Induction of the cytochrome oxidase system reduces the levels of a variety of other compounds alongside its autoinduction effect. The general principle is one every prescriber can apply without memorising a table: any co-prescribed drug cleared by the induced enzyme will fall in concentration over the same weeks, and the clinical presentation is loss of efficacy rather than toxicity. The full cross-class matrix of which psychotropics are affected belongs to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

Two consequences are properly part of this pair's pharmacology rather than of any interaction table. The first is within the mood stabiliser class itself: carbamazepine approximately halves lamotrigine levels, so a lamotrigine dose that is correct alongside carbamazepine is not the dose that would be correct alone. The titration schedule that follows from this is developed in the section of these notes on lamotrigine and the titration that prevents the rash. The second is contraceptive failure, which is a direct kinetic consequence of induction rather than a separate reproductive question: a patient established on carbamazepine who requires contraception should receive a preparation containing not less than 50 micrograms of oestrogen, or should use a non-hormonal method instead.

21.5 De-induction, or the same mechanism run backwards

Induction is a change in the amount of enzyme, so removing the inducer does not restore the previous state immediately; the enzyme population has to fall back by ordinary protein turnover. That reverse process, **de-induction**, is the half of the mechanism that examinations reward and clinics forget. When carbamazepine is stopped, for example on a switch to oxcarbazepine, de-induction of the CYP3A4 enzyme is to be expected, and substantial increases in lamotrigine levels would occur.

Notice the asymmetry in how the two directions feel to the prescriber. Starting an inducer produces falling concentrations and a slow loss of effect, which presents as a treatment that has stopped working and is easy to attribute to the illness. Stopping an inducer produces rising concentrations in every drug that was being cleared faster, which presents as new toxicity in a patient whose doses were not touched. The second is the more dangerous of the two, because nothing on the prescription changed except a drug that was taken away, and the temporal

association points in the wrong direction. Any manoeuvre involving carbamazepine, in either direction, changes the clearance of everything else the patient is taking.

TRAP **Treating the withdrawal of carbamazepine as a pharmacologically neutral act.** The stem gives a patient on carbamazepine and lamotrigine who is switched from carbamazepine to oxcarbazepine, and asks what to do about the lamotrigine. The expected wrong answer is nothing, because the lamotrigine dose has not been altered and the replacement drug looks like a near relative. The pharmacology says the lamotrigine dose was only ever correct in the presence of an inducer, that de-induction will raise its concentration substantially, and that the lamotrigine therefore needs attention at the moment the carbamazepine comes off. Candidates who have learned only that inducers lower levels answer half the question.

21.6 Oxcarbazepine: not a potent inducer, and not inert either

The derivative was designed around exactly the problem set out above, and on the central point it succeeds. Oxcarbazepine shows no autoinduction, and few major drug interactions have been reported. A clinician therefore inherits a drug whose clearance does not move underneath them, so there is no induction curve to time a level against and no drift in efficacy over the first weeks.

What the derivative does not do is escape the system altogether, and this is where the comparison earns its marks. Oxcarbazepine remains susceptible to induction by others: strong inducers of cytochrome P450 enzymes, such as carbamazepine, reduce plasma levels of its **monohydroxy derivative** metabolite by 29 percent to 40 percent. The contrast is one of degree and of direction: carbamazepine is a potent inducer and oxcarbazepine is not, and in an interaction oxcarbazepine is usually the affected drug rather than the inducing one. It has not shed inducing activity altogether, since oestrogen levels in oral contraceptives are still reduced by it.

AXIS OF COMPARISON	CARBAMAZEPINE	OXCARBAZEPINE
Principal action	Blocks voltage-dependent sodium channels, inhibiting repetitive neuronal firing	Blocks the same channels, and additionally increases potassium conductance and modulates high-voltage activated calcium channels
Active molecule	The parent compound itself	A pro-drug for licarbazepine
Effect on cytochrome P450	A potent inducer, particularly of CYP3A4	Not a potent inducer
Autoinduction	Present: an initial plasma half-life of around 30 hours falls to around 12 hours on chronic dosing	Absent
Role when inducers are involved	The inducer, of other compounds and of itself	The affected drug: strong inducers reduce the monohydroxy derivative metabolite by 29 percent to 40 percent
Hormonal contraception	For a patient needing contraception, a preparation containing not less than 50 micrograms of oestrogen, or a non-hormonal method	Oestrogen levels in oral contraceptives are still reduced

GOOD TO REMEMBER

Switching a patient from carbamazepine to oxcarbazepine solves the induction problem for the future and creates a de-induction problem for the present. Both have to be managed in the same consultation. Treat it as two separate pharmacological events: one drug leaving, which raises the concentration of everything it was clearing faster, and one drug arriving, which is stable in itself but still reduces contraceptive oestrogen and is still vulnerable to any other inducer. A switch documented as a straight substitution has recorded only half of what happened.

21.7 The rash, the allele and the ancestry question

The dermatological reaction is the other examinable axis of carbamazepine and it is entirely independent of enzyme induction. Around **3 percent** of patients treated with carbamazepine develop a **generalised erythematous rash**, and serious exfoliative reactions can rarely occur, with a vulnerability that is genetically determined. Those two statements describe different phenomena and should never be merged. The common rash is a frequent, generally benign event; the exfoliative reaction is rare, potentially catastrophic, and carries an identifiable genetic predisposition. Conflating them either terrifies clinicians out of a useful drug or, more dangerously, makes them casual about a rash that is not the common one.

The genetic predisposition is where the pharmacology becomes actionable before the first tablet. The **human leukocyte antigen (HLA) variant B*15:02** has a sensitivity of around 70 percent and a specificity approaching 100 percent in certain populations for serious skin reactions. Read

the two figures as the test's two different jobs. The very high specificity means few patients who are not susceptible test positive, so a positive result is a reason to choose another drug rather than a prediction that a reaction would have followed; with no predictive value and no population prevalence available, the post-test probability cannot be computed. The sensitivity, which is far from complete, means a negative result reduces risk without abolishing it, so the clinical vigilance that applies to every patient on this drug is not suspended by a reassuring test. On the strength of that performance, genetic testing is recommended in people from South-East Asia before carbamazepine is prescribed, a recommendation expressed by ancestry rather than by nationality or by current residence.

IN INDIA

Three decisions follow, and none of them depends on a figure this chapter does not have. Ask and record ancestry before writing the first carbamazepine prescription, because the testing recommendation issued in the United Kingdom guidance is framed by ancestry and can only be applied by a clinician who has asked the question. Where the ancestry brings the patient within that recommendation, test before prescribing rather than after a reaction. And counsel every patient started on the drug, tested or not, to stop it and present the same day if a rash appears, because a negative test still leaves a substantial minority of serious reactions undetected. The sources consulted here state no India-specific testing position and no allele frequency for Indian populations, so no such position is asserted; the ancestry recommendation is reported exactly as it is issued.

21.8 Marrow and sodium: the effects that owe nothing to enzymes

Two further adverse-effect groups complete the profile, and both are commonly misremembered because a rare event and a common one sit side by side in each. On the haematological side, carbamazepine commonly causes a **chronic low white cell count**, and one patient in 20,000 develops agranulocytosis or aplastic anaemia. These are not two grades of the same process. The benign leucopenia is frequent, usually stable and rarely a reason to stop; the marrow failure is rare and is a different event entirely. The examination distinction is the same one that separates the two rashes, and the practical distinction is that a modestly low count in a well patient is not the signal, whereas fever, mucosal ulceration or bruising in the same patient is.

On the electrolyte side, dry mouth, oedema and **hyponatraemia** are common with carbamazepine. Hyponatraemia is the one to carry forward, because it is common enough to be expected, easy to overlook when its symptoms are read as illness or as sedation, and additive with other exposures. The risk of hyponatraemia may be increased by other drugs that have the potential to deplete sodium, such as diuretics. That additivity is also what makes carbamazepine an indirect precipitant of lithium toxicity through sodium depletion; the toxic state itself, and the precipitants a prescriber controls, are developed in the section of these notes on lithium toxicity.

MNEMONIC

Self, Substrates, Stopping

Self: the drug induces its own metabolism, so its half-life falls and its own concentration drifts down over the first weeks. **S**ubstrates: the same induction lowers the concentration of other compounds cleared by the induced enzyme, which presents as loss of efficacy rather than as toxicity. **S**topping: de-induction reverses both, so withdrawal raises the concentration of everything that was being cleared faster. Oxcarbazepine deletes the first, blunts but does not abolish the second, since contraceptive oestrogen still falls, and stays susceptible to induction by others.

Carbamazepine induces CYP3A4 and is itself a substrate of it, so its initial plasma half-life of around 30 hours is reduced to around 12 hours on chronic dosing and the levels of co-prescribed substrates fall with it; oxcarbazepine shares the sodium channel action without the autoinduction, and stopping carbamazepine reverses both effects.

22 · Lamotrigine: the titration that prevents the rash

Almost every other drug in this chapter is titrated upward to find the dose that works; lamotrigine is titrated upward to avoid the dose that harms. That inversion is the organising fact of its pharmacology. The molecule is undramatic at the receptor, but the schedule by which it is introduced is a safety intervention with a named adverse outcome behind it, and it is one of the few dosing schedules a candidate must reproduce exactly. What follows develops the titration that prevents the rash as pharmacology rather than as a table to be memorised: how the body disposes of the drug, what two co-prescribed agents do to that disposal, and why every step of the schedule falls out of the kinetics.

Two boundaries are worth fixing before the detail. Lamotrigine as an antiepileptic, its titration in that setting and the rash risk as a problem of seizure treatment belong to the Epilepsy and psychiatry notes. Its place among the treatment options for bipolar disorder, including where it sits against lithium and valproate, belongs to the Mood disorders: pharmacotherapy notes. What is developed here is the systematic pharmacology both chapters assume: the disposition of the drug, the two interactions that dominate it, and the safety arithmetic that follows.

98 percent	55 percent	85 percent	10 percent
ORAL BIOAVAILABILITY	BOUND TO PLASMA PROTEIN	CONJUGATED TO INACTIVE GLUCURONIDES	EXCRETED UNCHANGED IN URINE

22.1 A mechanism that explains the effect and predicts none of the risk

The pharmacodynamic account is short. Lamotrigine **inhibits neuronal sodium channels**, and it may also inhibit glutamate release. Both are anticonvulsant actions, and the second is a plausible route to mood stabilisation given the excitatory load implicated in affective illness, but neither carries the confidence of the receptor-occupancy accounts of the antipsychotics. The

biophysics of channel gating and transmitter release belongs to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

The teaching point is what this mechanism does not do. Everywhere else in this chapter the adverse-effect signature is read forward from a binding profile: block this receptor, predict that complaint. Sodium-channel action generates no such map, and it is shared with other anticonvulsants used as mood stabilisers, so it discriminates nothing within the class. The adverse effect that governs how lamotrigine is prescribed is not a receptor effect at all: it is a rash, and no mechanistic detail about channel gating predicts it, times it or grades it.

That is why this section is shaped as it is. For an antipsychotic the receptor profile is the spine; for lamotrigine it is the disposition of the drug and the two interactions acting on it, because those set the exposure the dangerous adverse effect tracks.

22.2 A drug that is acted upon rather than one that acts

Lamotrigine is unusual among psychotropics in being almost entirely the victim of drug interactions rather than their author, and the reason is metabolic. Absorption is near complete, so the oral dose is effectively the systemic dose and no absorption variable to blame when an exposure surprises. Plasma protein binding is moderate, which keeps displacement out of the picture. The decisive fact is the route of elimination: the great majority of a dose undergoes **conjugation to inactive glucuronide metabolites**, and only a small fraction leaves the body unchanged in the urine.

Conjugation rather than oxidation is what makes this drug behave as it does in a polypharmacy regimen. The cytochrome enzymes that organise most psychotropic interaction risk are not the enzymes handling lamotrigine, which is why it is primarily a target of interactions rather than an instigator of them. What it does possess is a conjugation pathway other drugs can accelerate or block, and that asymmetry generates the rest of this section. The psychotropic substrate map is developed in the section of these notes on cytochrome P450 in practice, and the cross-class interaction matrix in the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

RULE **Lamotrigine is primarily a target of interactions rather than an instigator of them.** When a regimen containing lamotrigine misbehaves, the first question is not what lamotrigine has done to the other drug. It is what the other drug has done to lamotrigine's conjugation, and therefore to the exposure the titration was designed around. The interactions that matter here run in that one direction, and a candidate who has fixed it can reconstruct the consequences without memorising them.

22.3 One drug, three half-lives

The same tablet behaves as three different drugs depending on what else the patient is swallowing, and the variable is the elimination half-life. Taken alone, lamotrigine has a half-life of around a day and a modest clearance. Add a hepatic enzyme inducer such as carbamazepine or phenytoin and the half-life falls by about 50 percent while clearance doubles. Add valproate, which blocks the conjugation pathway, and the half-life doubles while clearance falls by

approximately 50 percent. The two interactions move the same parameter in opposite directions by comparable magnitudes, which is why neither can be treated as a minor adjustment.

Two consequences follow. The first is temporal: half-life sets how long an increment takes to declare itself, so an increase made in a patient on valproate is still climbing towards its plateau when the same increase in a patient on carbamazepine has long since settled. The second is exposure: a longer half-life at an unchanged dose means a higher steady concentration, and the steady concentration is what the serious adverse effect tracks. The general relation between half-life, steady state and the timing of a dose change is developed in the section of these notes on pharmacokinetics and half-life.

The three metabolic environments compare cleanly.

METABOLIC CONTEXT	LAMOTRIGINE HALF-LIFE	LAMOTRIGINE CLEARANCE	EFFECT ON LAMOTRIGINE CONCENTRATION
Lamotrigine alone	Around 28 hours	About 40 millilitres per minute	The reference against which the others are read
With an enzyme-inducing anticonvulsant	Approximately 14 hours	About 80 millilitres per minute	Decreased by 40 to 50 percent
With valproate	Approximately 56 hours	Around 20 millilitres per minute	Approximately doubled

22.4 Two directions, and the agents that produce them

The inducers that matter are carbamazepine, phenytoin, phenobarbital, primidone, rifampicin and oestrogen-containing hormonal contraceptives, and every one of them pulls the lamotrigine concentration down. Valproate is the single important agent pulling the other way, and it approximately doubles the concentration. The list is worth learning as a list, because the only feature its members share is an effect on the conjugation pathway.

The contraceptive entry deserves emphasis: it is the one the psychiatrist does not write. A woman established on lamotrigine who begins an oestrogen-containing contraceptive loses a substantial fraction of her exposure without any alteration to the psychiatric prescription, and one who stops it regains that fraction. Nothing has changed on the drug chart; the metabolic environment around it has. What it does not change is which column the titration is read from. The three schedules are selected by valproate and by an enzyme-inducing anticonvulsant; the prescribing information does not extend the doubled column to a contraceptive, which alters the exposure a given dose delivers and not the step the prescription opens at. Prescribing across pregnancy, lactation and the other special populations belongs to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

Enzyme induction as a class phenomenon is developed in the section of these notes on carbamazepine and oxcarbazepine, and valproate's own pharmacology in the section on valproate

as a mood stabiliser. What matters here is only the consequence for lamotrigine: the schedule is selected by the co-prescription, not by the diagnosis.

GOOD TO REMEMBER

Two questions decide which of the three lamotrigine schedules applies, and they belong before the first tablet rather than at the first review. Is the patient taking valproate? Is the patient taking an **enzyme-inducing anticonvulsant**? A third question is asked for a different reason, because it changes the exposure without changing the column: is the patient taking an oestrogen-containing contraceptive? All three are taken from the patient rather than from the referral letter, because a contraceptive is frequently absent from a psychiatric drug list. A fourth follows at every review: has any answer changed?

22.5 Two rashes, with very different frequencies and one shared instruction

The adverse effect defining this drug arrives in two forms. Benign rashes are common, occurring in as many as **10 percent** of patients taking lamotrigine. Serious rashes are rare, and they are why the prescribing information carries a **boxed warning** naming Stevens-Johnson syndrome among the reactions requiring admission and withdrawal of the drug. The incidence of serious rash was estimated at 1 percent in prospectively monitored patients under 16 years and **0.3 percent** in premarketing clinical trials of adults. That estimate dates from 1997 and should be quoted as a product-information figure rather than as a current incidence.

What matters more is the shape of the risk than its magnitude, and that shape produces an instruction that looks disproportionate until the reasoning is visible. The prescribing information does not ask the clinician to distinguish the benign rash from the dangerous one; it advises that lamotrigine ordinarily be stopped at the first sign of rash unless the rash is clearly not drug related. That rule is deliberately blunt. It accepts that most drugs stopped under it were causing nothing worse than a nuisance, because the cost of that error is a lost treatment option and the cost of the opposite error is a patient in hospital.

Nearly all life-threatening rashes have occurred between **2 and 8 weeks** after lamotrigine is started, although isolated cases have been reported after prolonged treatment. The window is the practical content of that observation: it says when vigilance must be highest, and it coincides with the period the escalation occupies. The tail is equally examinable, because a rash appearing later is not exonerated by its timing.

22.6 The risk factors a schedule can remove

Lamotrigine remains in routine use because most of the identified risk is modifiable, and the modifications are properties of how it is introduced. Besides young age, the factors raising the risk of a serious rash are increasing the lamotrigine dose too quickly, concurrent valproate, which doubles lamotrigine concentrations, and exceeding the recommended initial dose or the recommended rate of escalation.

- Young age, which the prescriber does not control and which the two incidence figures already separate out.

- Increasing the dose too quickly, which the prescriber controls completely and is the commonest way the risk is manufactured.
- Concurrent valproate, which the prescriber controls at the level of the schedule even when the co-prescription is clinically necessary.
- Exceeding the recommended starting dose or escalation rate, the same error from the other end, and the one that arrives through a prescription written from memory.

Three of the four are features of the prescription rather than of the patient, and one instrument addresses all three. That is the sense in which the titration is a safety intervention rather than a dosing convention: the deliberate removal of three named risk factors.

22.7 The schedule, and the arithmetic behind it

The prescribing information sets out three schedules, far easier to hold as one schedule with two transformations than as three lists of numbers. **The reference schedule** is the one used when the patient is taking neither valproate nor an enzyme-inducing anticonvulsant; the other two columns are read against it.

TITRATION STEP IN BIPOLAR I MAINTENANCE (PRESCRIBING OWNED BY THE MOOD DISORDERS: PHARMACOTHERAPY NOTES)	WITH VALPROATE	WITH NEITHER VALPROATE NOR AN ENZYME-INDUCING ANTICONVULSANT	WITH AN ENZYME-INDUCING ANTICONVULSANT
Opening dose (first and second weeks)	12.5 mg daily	25 mg daily, the reference step, prescribed elsewhere	50 mg daily
First increase (third and fourth weeks), indication owned elsewhere	25 mg daily	50 mg daily, a step another chapter reaches for	100 mg daily
Second increase (fifth week), not taught here	50 mg daily	100 mg daily, prescribing owned by those notes	200 mg daily
Third increase (sixth week), another chapter's decision	100 mg daily	200 mg daily, a level chosen elsewhere	300 mg daily
Target dose (seventh week onward), owned by that chapter	100 mg daily	200 mg daily, a destination not taught here	400 mg daily

The arithmetic is nearly exact, and it is the examinable part. Every step in the valproate column is one half of the corresponding step in the reference column, without exception. The enzyme-inducer column is double the reference step at every stage but the sixth week, where the reference column has already reached its target and holds while the inducer column still has one increase to

make. Nothing else changes: the step boundaries fall in the same places and the escalation takes the same number of stages in each column.

The reason is the concentration rather than the milligram. Valproate approximately doubles the concentration produced by any given dose, so halving the dose returns the patient to the exposure the reference schedule was built around. An enzyme inducer moves the concentration a comparable distance in the opposite direction, so doubling the dose does the same work. The three columns are not three levels of caution but three prescriptions calculated to deliver one exposure profile.

MNEMONIC

Halve, Double, Restart

Halve every step when valproate is in the regimen, because valproate has already doubled the exposure. **D**ouble the opening step and the target when an enzyme-inducing anticonvulsant is in the regimen, because the inducer has removed a comparable share of it. **R**estart from the opening step after any meaningful interruption, because the exposure the schedule built has gone. Three words carry the prescription once the reference column is known, with the sixth week of the inducer column read off the table rather than off the rule.

TRAP **Slowing the titration for a patient on valproate while keeping the ordinary target dose.** The stem describes a patient established on valproate in whom lamotrigine is being added, and the expected wrong answer halves the early steps and then climbs to the usual maintenance figure over a few extra weeks. It is wrong because the transformation applies to the destination as well as to the journey: the valproate column ends lower, not merely later. A doubled concentration at the ordinary target dose is the very exposure the schedule exists to prevent.

22.8 Doses that come from trials, not from a level

It is tempting to read the three columns as three attempts to hit one blood concentration, and to conclude that a lamotrigine level would settle the matter. It would not. The prescribing information derives its target doses from controlled trials rather than from a plasma concentration target, and no therapeutic range has been established for lamotrigine in bipolar disorder. The **target dose** is therefore a trial result, not a surrogate for an unstated concentration.

The consequence inverts the assumption that a level is always the more rigorous instrument. Where lithium offers a number reporting what the patient is actually exposed to, lamotrigine offers only the schedule and the dose, so everything protecting the patient is front-loaded into the prescription: the co-prescription question, the escalation rate and the response to a rash. No later measurement catches an error made at the outset. The indications for measuring a psychotropic concentration, and the published ranges for the drugs that have them, belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes; the logic of a narrow therapeutic index is developed in the section of these notes on lithium as a narrow-index drug.

22.9 An interruption resets the schedule

The last piece of the arithmetic is the one most often lost, because it is triggered by an event nobody prescribed. If a patient does not take lamotrigine for **5 days** or more, which is about five times its half-life, the drug must be reintroduced using **the full gradual escalation regimen**, because serious rashes can occur when it is restarted rapidly.

The reasoning is entirely kinetic. After roughly five half-lives a drug has effectively left the body, so a patient resuming at the maintenance dose is not continuing a titration; they are beginning one at its endpoint. Whatever the gradual introduction achieved is not a property of the prescription that can be reinstated by writing it again. The interruption is common and rarely announced.

PITFALL

Restarting a patient at the dose written on the last prescription after a gap in supply. The clinician sees a stable patient on an established maintenance dose, the pharmacy dispenses that dose, and nobody records how long the gap actually was. The question that prevents this is trivial and almost never asked: how many days did the patient go without a tablet? A gap that crosses the threshold converts a repeat prescription into a fresh initiation, and it has to be written as one.

A drug whose elimination can be doubled or halved by a co-prescription, whose dangerous adverse effect tracks exposure and appears in a defined early window, and for which no concentration can be measured, has one place left to put its safety: the schedule.

Lamotrigine's titration is a safety intervention rather than a convention: halve every step and the target with valproate, double the opening step and the target with an enzyme-inducing anticonvulsant, stop the drug at the first sign of any rash unless it is clearly not drug related, and after an interruption of several days begin the whole escalation again.

Anxiolytics, hypnotics and other agents

23 · Benzodiazepines: tolerance, dependence and withdrawal as pharmacology

Almost everything clinicians disagree about with the benzodiazepines is settled at the receptor. Why the class is so forgiving taken alone and so dangerous taken in company, why one member sedates while another mainly relaxes muscle, and why the last part of a reduction is harder than the first, are all consequences of one receptor complex and of what a benzodiazepine molecule does to it. This topic sets out that pharmacology and then reads tolerance, dependence and withdrawal off it, as receptor-level phenomena with a mechanism rather than as failures of prescribing or of character.

The ground is examined heavily because it is one of the few places in psychopharmacology where a mechanism, a within-class distinction and a well-quantified clinical consequence sit in a straight line, so one question can test all three. The use of these drugs as treatments for anxiety and for

insomnia belongs to the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes; benzodiazepine dependence as a diagnosable substance use disorder, with its assessment and its detoxification, belongs to the Substance use disorders notes; and the biochemistry of GABA as a transmitter belongs to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes. What is owned here is the class as a pharmacological object.

58 to 100%	**15 to 44%**	**up to 15%**	**alpha-1**
OF LONG-TERM USERS HAVE A WITHDRAWAL REACTION	HAVE MODERATE TO SEVERE WITHDRAWAL ON STOPPING	DEVELOP A PROTRACTED WITHDRAWAL SYNDROME	THE SUBUNIT CARRYING SEDATION AND ATAXIA

23.1 The receptor, and what the drug actually does to it

The target is the GABA-A receptor, a ligand-gated chloride channel. It is a **pentamer**, and the assembly most commonly found in the brain is built from **two alpha subunits, two beta subunits and one gamma subunit**, which is the common arrangement rather than the only one. Benzodiazepines do not bind where GABA binds. They bind at **the interface of the gamma-2 and alpha subunits**, a site of their own, and that structural separation is the reason the class can only ever modify a signal that is already arriving.

The structural point becomes the defining pharmacological one. A benzodiazepine binds directly to a specific modulatory site on the receptor and enhances the effect of GABA, but it does not open the chloride channel in the absence of GABA. It is therefore a **positive allosteric modulator** and not a direct agonist. That is not hair-splitting, it is the origin of the class's safety profile: the drug supplies no inhibition of its own, it only amplifies inhibition the brain is already generating. What the allosteric change does is increase the **frequency with which the central chloride channel opens**. More openings per unit time, on a channel that will not open at all unless the transmitter is present.

The barbiturates are the instructive contrast, and it is architectural rather than a matter of strength. They act at a different site on the same complex, and at higher concentrations they become direct GABA-A agonists. A drug that can open the channel unaided has no ceiling written into its mechanism; a drug that can only amplify an existing signal has one built in before the first prescription. Two groups producing a similar gross clinical effect are therefore not pharmacologically interchangeable, and the question is settled at the binding site.

RULE **A benzodiazepine modulates; it does not open.** Every distinctive property of the class descends from that one sentence. The safety in isolated overdose follows from it, because the effect is bounded by how much GABA the brain is releasing. So does the shape of withdrawal, because a system that has adapted to chronic amplification of its own inhibition is left under-inhibited when the amplifier is removed.

TRAP **Writing that a benzodiazepine increases the duration of chloride channel opening.** The action is on the *frequency* of opening, and frequency is a distinct parameter from duration. Candidates memorise the benzodiazepine and barbiturate pair as a matched contrast, then reconstruct it under time pressure from whichever half they recall first. Learn the benzodiazepine half on its own terms, as frequency of opening at a modulatory site that requires GABA to be present, and the answer stops depending on the other half.

23.2 Subunit selectivity, and why one class has several personalities

Not every GABA-A receptor responds to a benzodiazepine, and the discriminating component is the alpha subunit, of which there are six numbered subtypes. Receptors carrying alpha-1, alpha-2, alpha-3 or alpha-5 are **benzodiazepine-sensitive**, while the presence of **alpha-4 or alpha-6** leads to a lack of sensitivity to the drugs altogether. A fraction of the inhibitory machinery is therefore invisible to the class, worth holding whenever an effect is attributed to global GABA enhancement.

Among the sensitive receptors, the subunit predicts which effect is produced. Receptors containing alpha-1 appear to mediate the sedative and hypnotic effect, and the same subunit is implicated in ataxia. The antianxiety effect is attributed instead to receptors containing **alpha-2 and possibly alpha-3**, and those two subunits also carry the muscle relaxant effect. The hedging is the source's own and should be preserved: the alpha-3 attribution is stated as a possibility, not a finding.

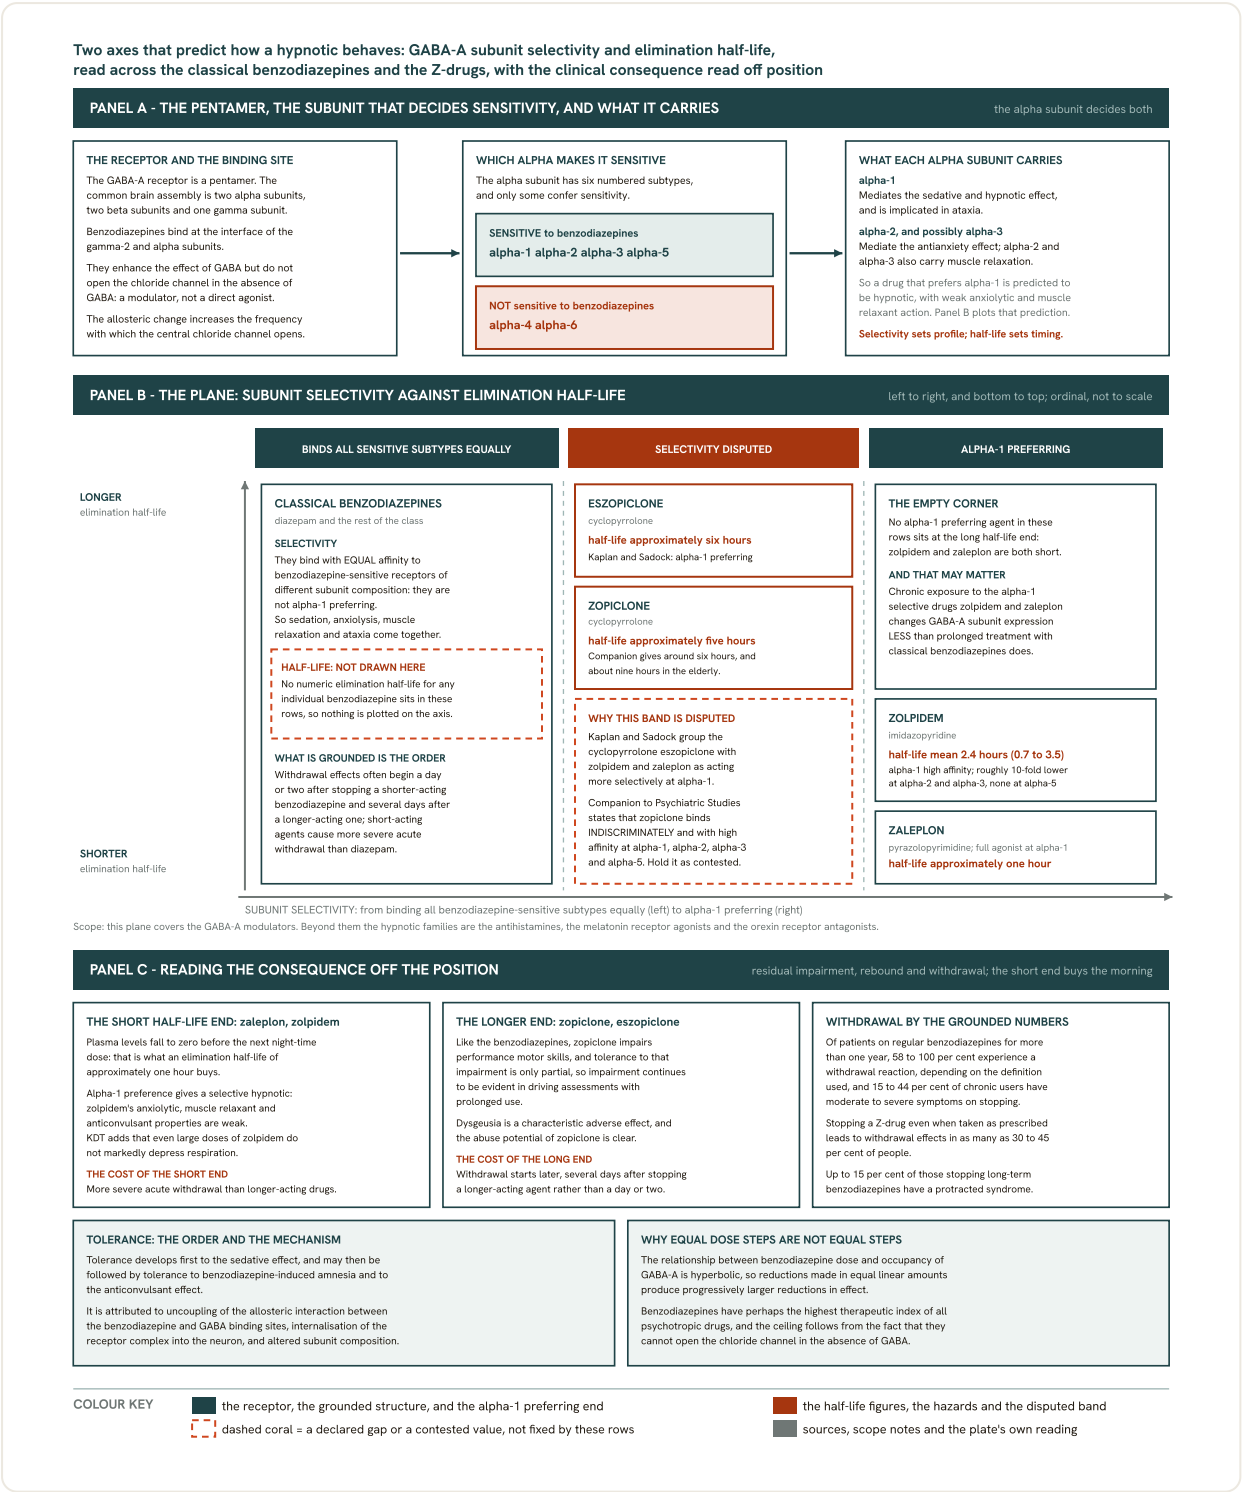

Fig 36.9 GABA-A subunit selectivity set against elimination half-life. Subunit preference predicts the profile and half-life predicts the timing: alpha-1 carries sedation, hypnosis and ataxia, while alpha-2 and possibly alpha-3 carry anxiolysis and muscle relaxation. The classical benzodiazepines bind benzodiazepine-sensitive receptors of different subunit composition with equal affinity and are drawn without a half-life, because no numeric half-life for an individual benzodiazepine sits in the rows behind this plate; the Z-drugs are plotted in half-life order, with the cyclopyrrolones marked as the contested selectivity case. The short end buys a clear morning and pays in more severe acute withdrawal, the longer end carries residual motor impairment and a later withdrawal onset. (Kaplan and Sadock, CTP 2024; Companion to Psychiatric Studies; The Maudsley Deprescribing Guidelines; KDT Essentials of Medical Pharmacology.)

ALPHA SUBUNIT PRESENT	BENZODIAZEPINE SENSITIVITY	EFFECT ATTRIBUTED TO IT
alpha-1	Sensitive	Sedative and hypnotic effect; implicated in ataxia
alpha-2	Sensitive	Antianxiety effect, and muscle relaxation
alpha-3	Sensitive	Possibly antianxiety; also muscle relaxation
alpha-5	Sensitive	No effect attributed to it in these sources
alpha-4 and alpha-6	Not sensitive	No benzodiazepine action at these receptors

Two consequences follow, and the accompanying figure lays the same logic out alongside the kinetic axis. The first is pharmacological ambition: if sedation and anxiolysis are carried by different subunits, a subunit-selective molecule could in principle separate them, which is the reasoning behind every attempt to build an anxiolytic that does not sedate. The second is diagnostic. When one patient is unsteady and another on a different member is chiefly calm, that is not idiosyncrasy, it is a difference in where the molecule works laid over a difference in how long it stays there. Half-life and steady state have their own section in these notes, as does the comparison with the Z-drugs, which exploit precisely this subunit axis.

23.3 The therapeutic index, and the one route to respiratory harm

The mechanical ceiling has a measurable consequence. Benzodiazepines have perhaps the highest **therapeutic index** of all psychotropic drugs, a remarkable position for a class this sedating to occupy. It is the arithmetic of a modulator: escalating the dose recruits more amplification of an inhibitory signal that is itself finite, so the curve flattens where an agonist's would keep climbing. What a narrow index means in general, and which psychotropics consequently need a concentration measured, has its own section in these notes.

The ceiling does not abolish respiratory risk; it relocates it. At therapeutic doses the class does not interfere with respiration, but rapid intravenous administration of a swiftly distributed compound such as diazepam can produce significant respiratory depression, and this is reversible with the benzodiazepine antagonist **flumazenil**. Route and speed, not dose alone, generate the hazard, and because the drug occupies a defined site another molecule can displace it from that site, a reversibility no metabolic strategy could deliver.

The historical comparison finishes the argument. Barbiturates are blighted by high dependency potential and a very narrow therapeutic index, and the same problems afflicted their analogues glutethimide and methaqualone, popular in the 1960s and 1970s. The benzodiazepines displaced them not by being better sedatives but by having a mechanism with a ceiling, and the property that made the class safe is also what made it prescribable at the scale that created the dependence problem.

GOOD TO REMEMBER

The high therapeutic index is a statement about the drug alone. It does not survive combination with other central nervous system depressants, because the ceiling is on this drug's modulation of GABA, not on total central depression from every source acting at once. A patient told the tablets are very safe has been told something true about a situation that does not describe them if they also drink, or if another prescriber adds a second sedating agent. The sentence that carries the pharmacology accurately is that the drug is safe by itself, conditional on staying by itself. Interactions across drug classes are developed in the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

23.4 Tolerance: the order it arrives in, and what moves at the receptor

Tolerance to this class is not a single event but a sequence, and the sequence is examinable. It develops to many of the drugs' actions but not all, and it develops first to **the sedative effect**. That may then be followed by tolerance to benzodiazepine-induced amnesia and to the anticonvulsant effect. The clinical reading is immediate: the effect a patient most notices at the start is the effect most likely to disappear, which is the pattern that produces dose escalation in someone who is not drug-seeking but simply chasing back an effect that has faded.

MNEMONIC

Sedation first, then amnesia, then anticonvulsant

Sedation goes first and goes fastest, which is why the initial sleepiness is a poor guide to how the drug will feel later. **A**mnesia may follow. **A**nticonvulsant effect may follow as well, and that is the limb with the sharpest consequence, because an agent relied on for seizure control cannot be assumed to hold that effect indefinitely. The order runs from the effect the patient reports to the effect nobody notices until it is needed.

Underneath the sequence there is a receptor-level account, and the honest version is that the sources call the mechanism unclear and describe it as linked to dose and duration of exposure. What is named is a set of adaptations:

- **Uncoupling** of the allosteric interaction between the benzodiazepine and GABA binding sites, so the drug's occupancy stops translating into enhancement.
- Internalisation of the GABA-A receptor complex into the neuron, which removes receptor from the membrane altogether.
- Alteration in the subunit composition of the receptors, which changes what the surviving population responds to.

Learn them together: they explain why tolerance is not reversed by a dose increase, and why the receptor a patient has after long exposure is not the one they started with. Receptor up-regulation, down-regulation and sensitisation as general phenomena belong to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

23.5 Withdrawal is the ordinary case, not the exception

The single most common misconception about this class is that withdrawal is what happens to a minority who misuse it. The figures say otherwise, and they are worth carrying away. In patients who have taken regular benzodiazepines for more than one year, 58 to 100 per cent experience a withdrawal reaction, the width of that band reflecting how differently studies define the event, not any doubt that it occurs. Narrowing to clinically substantial symptoms, **15 to 44 per cent** of chronic users experience moderate to severe withdrawal symptoms on stopping.

Read those two numbers as one statement with two thresholds. Some degree of withdrawal after a year of regular exposure is the expected pharmacological outcome, while a substantial minority have symptoms severe enough to dominate the picture. Neither figure describes a substance use disorder. They describe physiological adaptation in patients taking the drug as prescribed, and conflating the two attributes to the patient's conduct what belongs to the receptor. The disorder, its diagnosis and its treatment belong to the Substance use disorders notes.

> **PITFALL**
>
> Reading withdrawal as relapse. When anxiety, insomnia and autonomic arousal appear during or after a reduction, they overlap almost entirely with the presentation the drug was started for, so the natural inference is that the underlying disorder has returned and the dose should go back up. Given how commonly a withdrawal reaction follows long-term use, the prior probability sits the other way. The discriminators are temporal and pharmacological rather than phenomenological: the relationship of onset to the reduction, and the half-life of the agent reduced. Getting this wrong does not merely delay a taper, it teaches the patient that they cannot manage without the drug.

23.6 The timetable is written by half-life, but not the whole of it

Onset of withdrawal tracks elimination, the most directly useful kinetic prediction here:

- After a shorter-acting benzodiazepine, effects often begin a day or two after stopping.
- After a longer-acting one they often begin several days after stopping, and delayed-onset effects are also recognised.

Severity follows the same axis in the acute phase. Short-acting benzodiazepines are associated with more severe acute withdrawal than longer-acting drugs such as diazepam, which is what the pharmacology predicts: a concentration falling steeply gives the adapted system no time to readjust, while a slow decline is a taper the drug performs on the patient's behalf. That principle explains most of the within-class variation clinicians encounter, and it is the content behind every equivalence table, since potency and half-life vary independently and a table is needed precisely because knowing one tells you little about the other. The numeric equivalences are approximations rather than conversions, and belong to the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes.

What half-life does not predict is the long tail. **A protracted withdrawal syndrome** is estimated to occur in up to 15 per cent of patients stopping long-term benzodiazepines, lasting months and

sometimes years, and both short and long acting agents can produce it. Regulation has followed the evidence: in 2020 the FDA boxed warning on marketed benzodiazepines was updated to state that withdrawal syndromes could persist for **more than 12 months** after stopping.

TRAP **Concluding that a long half-life protects against withdrawal.** It softens the acute phase, which is real and useful, and it does nothing for the protracted syndrome, which both short and long acting agents can produce. Candidates who compress the two findings into a single rank order answer confidently and wrongly. Hold them apart as two questions: how severe the acute phase is, which half-life governs, and whether a syndrome persists for many months, which half-life does not govern.

23.7 Why equal dose steps are not equal effect steps

The last piece of mechanism is the one most often missing from a candidate's account, and it explains why the end of a reduction behaves differently from the start. The relationship between benzodiazepine dose and occupancy of GABA-A is **hyperbolic**, not linear. Occupancy rises steeply at low doses and then flattens, so most of the receptor effect has already been bought by a small fraction of a typical dose.

Run the curve backwards and the clinical implication appears. Reductions made in equal linear amounts produce progressively larger reductions in effect, because each identical decrement moves further down the steep part of the curve than the one before it. A patient who tolerated the early reductions and then struggles badly on the last ones is not deteriorating or failing; they are meeting the geometry of the curve, which the pharmacology predicted before the reduction began. The corollary is that a plan reasoned in equal steps of milligrams is reasoning in the wrong currency, because the currency that matters is receptor occupancy. The schedules that follow belong to the prescribing and deprescribing lens rather than to class pharmacology, and are found in the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes.

23.8 The edges of what class pharmacology settles

Two limits deserve stating plainly, because both invite a manufactured answer. The first is that this account is built on classical benzodiazepines acting at benzodiazepine-sensitive receptors, and nothing above transfers automatically to a molecule that binds elsewhere. The second is jurisdictional: these sources record a United States regulatory warning and say nothing about scheduling in India, and a scheduling status recalled from memory is exactly the kind of confident specific that turns out to be wrong.

IN INDIA

The agent Indian practice must reason about, and that the standard textbook chapters do not cover, is etizolam, a thienotriazolodiazepine that behaves as a full benzodiazepine receptor agonist. The prescribing decision that follows is to treat it as a full member of this class for every purpose set out above: the same modulatory mechanism, the same subunit-determined split between sedation and anxiolysis, tolerance developing first to sedation, and a withdrawal reaction after long regular use. What must not be done is to borrow a half-life from another member to fill the gap, because the sources used here print none for it, and an assumed half-life silently decides how a reduction is planned. Its regulatory position needs the same discipline: read the current classification and the record-keeping it requires off the Drugs and Cosmetics Rules themselves rather than off a textbook.

What remains is compact and durable: one modulatory site with a ceiling built into the mechanism, a subunit map assigning sedation, anxiolysis and muscle relaxation to different parts of the same complex, a tolerance sequence beginning with the effect the patient notices, a withdrawal reaction that is the ordinary consequence of long exposure, and a dose-occupancy curve explaining why reductions get harder rather than easier.

A benzodiazepine cannot open the chloride channel without GABA, and the rest of the class follows from that: a ceiling that gives it perhaps the highest therapeutic index of all psychotropic drugs, a subunit map that splits sedation from anxiolysis, tolerance that takes sedation first, and a withdrawal reaction most long-term users will experience.

24 · Z-drugs and the other hypnotics compared

A hypnotic is compared with another hypnotic on two pharmacological axes and very little else. The first is which population of receptors the molecule prefers; the second is how long it persists in the plasma. Almost every question asked about this group is derived from those two axes: why one agent leaves impairment behind in the morning and another does not, why a drug with a weak anticonvulsant action is nonetheless an effective hypnotic, and why the withdrawal picture is so familiar. The **benzodiazepine receptor agonists** are the group in which the two axes separate most cleanly, which is why they repay being learned as pharmacology rather than as a list of names.

Two warnings about the label, because both are examined. The prefix in "non-benzodiazepine hypnotic" is a statement about chemical structure, not about mechanism and not about safety. And the letter Z is an orthographic accident: it groups agents that share a first letter and a binding site while differing in chemistry, in subunit preference and in the length of time they stay in the body. The GABA-A receptor itself, its subunit architecture and the biochemistry of inhibitory transmission belong to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes and are assumed here. Benzodiazepine class pharmacology has its own section in these notes; the use of any of these agents as a clinical treatment for insomnia

belongs to the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes, and sedative-hypnotic dependence as a disorder to the Substance use disorders notes.

about 1 hour	mean 2.4 hours	5 to 6 hours	about 6 hours
ZALEPLON, ELIMINATION HALF-LIFE	ZOLPIDEM, IMMEDIATE RELEASE	ZOPICLONE, ADULT FIGURE	ESZOPICLONE

24.1 Three chemistries, one binding site

The group is named for a letter and not for a structure. Zolpidem is an **imidazopyridine**, zaleplon a **pyrazolopyrimidine** and eszopiclone a **cyclopyrrolone**: unrelated chemical families that converge on the same pharmacological place. What unites them, on the account given in Kaplan and Sadock's Comprehensive Textbook of Psychiatry, is that they act more selectively than the classical benzodiazepines on GABA-A receptors containing the alpha-1 subunit, receptors often referred to as GABA-A1 receptors. The comparison is the load-bearing half of that sentence, and the half candidates drop. Diazepam and the other classical benzodiazepines bind with equal affinity to benzodiazepine-sensitive receptors of differing subunit composition; the Z-drugs prefer one population over the others.

That one comparative statement generates most of what follows. If a classical benzodiazepine engages every benzodiazepine-sensitive receptor more or less equally, its clinical profile will be the sum of everything those receptors do, and the adaptive consequences of occupying all of them for months will be equally broad. A molecule that prefers a single population should, on the same reasoning, deliver a narrower set of effects and adapt a narrower set of receptors. The Z-drugs are the test of that prediction, and the instructive finding is that it is borne out only in part.

RULE **The prefix records chemistry; the binding site records pharmacology.** Every drug in this group is an agonist at the benzodiazepine recognition site of the GABA-A receptor. Everything that follows from occupying that site therefore remains available: sedation, psychomotor impairment, tolerance and withdrawal on stopping. What differs between the agents is which receptor populations they prefer, and for how long they stay. Nothing in the word "non-benzodiazepine" suspends the first list.

24.2 Selectivity is graded, and it belongs to the molecule

Zolpidem is the cleanest instance. Companion to Psychiatric Studies describes it as binding selectively and with high affinity at alpha-1 containing GABA-A receptors, with roughly ten-fold lower affinity for alpha-2 and alpha-3, and none at all for alpha-5. Zaleplon sits beside it: a **full agonist** at the benzodiazepine alpha-1 receptor, which it binds with high affinity, its affinity for the alpha-2 and alpha-3 subtypes being lower. Different molecules, unrelated chemical families, one preference.

The clinical translation of that binding pattern is stated by the same source and is the part to carry away: the anxiolytic, muscle relaxant and anticonvulsant properties of zolpidem are weak. KDT Essentials of Medical Pharmacology, the Indian standard pharmacology text, makes the same attribution from the opposite direction, ascribing the relatively selective hypnotic effect of zolpidem to its preferential action at the alpha-1 subunit containing benzodiazepine receptor,

with no evident anticonvulsant, muscle relaxant or antianxiety effect to accompany it and an elimination half-life of about two hours. Independent texts, opposite framings, the same pharmacological claim.

Put those statements side by side and the mechanism becomes visible without anyone having to assert a subunit-by-subunit map of function that neither source spells out here. The actions a selective agent loses are carried by receptor populations it barely engages; the action it keeps is carried by the population it prefers. That is the entire promise of subunit selectivity, and it is why these molecules were developed.

Two riders keep the account honest, and both are examinable. Selectivity is not exclusivity: despite its preference, zolpidem is still displaced from its binding sites by **flumazenil**, which is the pharmacological proof that it is sitting where a benzodiazepine sits. And the narrowing of effect is not a narrowing of every risk. KDT records that even large doses of zolpidem do not markedly depress respiration, which is a genuine and useful contrast with the older sedatives, but respiratory margin is one property among many and it says nothing about the impairment, tolerance and withdrawal discussed below.

24.3 The cyclopyrrolones, where the tidy account breaks

The neat story survives exactly as far as the cyclopyrrolones. Companion to Psychiatric Studies states that while zopiclone is chemically a cyclopyrrolone, pharmacologically it resembles the benzodiazepines in binding indiscriminately and with high affinity across the benzodiazepine site, at alpha-1, alpha-2, alpha-3 and alpha-5 containing receptors. Kaplan and Sadock's account, quoted earlier, groups eszopiclone with zolpidem and zaleplon among the agents acting more selectively at the alpha-1 population. Both statements are in print in standard texts. They cannot be reconciled by rounding, and neither is a misprint.

The defensible reading is that selectivity here is a graded property rather than a binary one, and that it is strongest and least disputed for zolpidem and zaleplon. The cyclopyrrolones are the contested case, and a candidate is safest holding them that way: as the agents whose selectivity is claimed by one authority and denied by another, rather than as agents whose selectivity is settled. It is also worth noticing that the two statements are made about two different molecules, zopiclone in one text and eszopiclone in the other. Neither source, as quoted here, states how those two molecules are related to each other, so the reader should not assume that a finding about one transfers automatically to the other.

TRAP **Treating preference for one subunit population as the definition of a Z-drug.** Candidates learn the group as "the selective ones" and then meet a question about the cyclopyrrolones, where one standard text groups the class with the selective agents and another describes indiscriminate binding across the whole benzodiazepine site. Whichever way the paper is written, the memorised equation answers it wrongly. Learn the axis, not the slogan: selectivity is graded, it is strongest for zolpidem and zaleplon, and the cyclopyrrolones are where the sources disagree.

24.4 The second axis: how long the drug is there

Half-life does more comparative work in this group than receptor preference does. The **elimination half-life** of these agents spans a wide range for drugs given at the same point in the evening for the same purpose, and that range is what decides how much drug is still present when the patient has to function. The Maudsley Deprescribing Guidelines set the figures out drug by drug. Zaleplon is the extreme case, with an elimination half-life of approximately one hour and peak plasma concentration reached approximately one hour after oral ingestion, so that plasma levels fall to zero well before the next night-time dose is due. Zolpidem has a short but appreciably longer profile: a **mean of 2.4 hours** with a range of 0.7 to 3.5 hours, a duration of action of up to 6 hours, and peak plasma reached between 0.5 and 3 hours. The controlled release formulation produces a longer period of maximal dose and a half-life of 2.8 hours.

The cyclopyrrolones sit at the long end. Eszopiclone has an elimination half-life of **approximately 6 hours**, with peak plasma at approximately one hour. Zopiclone is the one figure the sources dispute. The Maudsley Deprescribing Guidelines give **approximately 5 hours**, with peak plasma at approximately 1.5 to 2 hours; Companion to Psychiatric Studies gives around six hours. The sensible teaching figure for an adult is therefore five to six hours, quoted with both sources named rather than resolved by picking one. The same text adds that the half-life is considerably longer in older people, at about nine hours, and that figure should be carried separately as an age-qualified statement rather than folded into the adult range. Prescribing in the elderly is treated in the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

AGENT	CHEMICAL CLASS	BINDING AT THE BENZODIAZEPINE SITE	PEAK PLASMA	ELIMINATION HALF-LIFE
Zaleplon	Pyrazolopyrimidine	Full agonist at alpha-1, high affinity; lower at alpha-2 and alpha-3	About one hour	About one hour
Zolpidem	Imidazopyridine	High affinity at alpha-1, roughly ten-fold lower at alpha-2 and alpha-3, none at alpha-5	Between 0.5 and 3 hours	Mean 2.4 hours, range 0.7 to 3.5 hours
Zopiclone	Cyclopyrrolone	Indiscriminate across alpha-1, alpha-2, alpha-3 and alpha-5, on the Companion account	About 1.5 to 2 hours	About five to six hours, sources differing
Eszopiclone	Cyclopyrrolone	Grouped with the alpha-1 preferring agents by Kaplan and Sadock	About one hour	About 6 hours

PITFALL

Reading a mean half-life as though it were the patient's own. The zolpidem figure is a mean drawn from a wide range, and the drug's stated duration of action is longer than that mean would suggest. A patient at the slow end therefore carries more drug into the morning than the textbook number predicts, and the grogginess gets attributed to poor sleep, to the illness or to anything except the tablet, on the reasoning that the half-life is short. Ask about the morning after, not only about the night.

24.5 What preference for one subunit population does not buy

If subunit selectivity delivered everything hoped of it, these agents would produce hypnosis without the residual burden of a benzodiazepine. They do not, and the evidence is at its clearest for zopiclone. Companion to Psychiatric Studies records that, like the benzodiazepines, it impairs performance motor skills; that tolerance to that impairment is only partial, so impairment continues to be evident in driving assessments with prolonged use; and that **dysgeusia**, a bitter metallic alteration in taste, is a characteristic adverse effect. The same source states that the abuse potential of the drug is clear.

The partial tolerance finding is the one that changes practice. Tolerance to sedation developing does not mean tolerance to impairment developing, and a patient who no longer feels drugged is not thereby a patient who performs normally. Abuse liability, dependence and their assessment as a disorder belong to the Substance use disorders notes; what belongs here is the pharmacological point that neither the impairment nor the abuse liability was abolished by moving away from benzodiazepine chemistry.

GOOD TO REMEMBER

Ask about driving at every review of a patient taking one of these agents, not only when the drug is started. The evidence that impairment persists comes from driving assessments in people on prolonged treatment, which is precisely the group a clinician is least likely to re-examine, because they are stable, they are not complaining, and the prescription has become a repeat. Habituation to the feeling of a drug is not evidence of recovered performance.

24.6 Receptor adaptation, and why stopping is still difficult

Here the selectivity argument does earn something measurable. Kaplan and Sadock's Comprehensive Textbook of Psychiatry reports that chronic exposure to the selective agents zolpidem and zaleplon produces less extensive change in the expression of GABA-A receptor subunits than prolonged treatment with classical benzodiazepines. That is a mechanistic result, not a marketing claim, and it is what the selectivity hypothesis predicts: fewer receptor populations occupied for months means fewer receptor populations remodelled. The important qualification travels with it. On short-term withdrawal, the changes that do occur come to resemble those seen after diazepam. Less adaptation during exposure, in other words, does not translate into a different kind of adaptation once the drug is removed.

The clinical arithmetic follows. The Maudsley Deprescribing Guidelines state that stopping these drugs, even when they have been taken exactly as prescribed, can lead to **withdrawal effects** in as many as **30 to 45 per cent** of people, and that the longer the drug has been taken and the higher the dose, the greater the risk. Two features of that statement deserve emphasis. It applies to prescribed use, so it is not a statement about misuse. And the same figure is given for each of the agents in the group, so it is not a peculiarity of one molecule that a switch within the class would solve.

TRAP **Reading "non-benzodiazepine" as "not dependence forming", and stopping abruptly on that basis.** The withdrawal frequency quoted above is not a benzodiazepine figure borrowed for effect; it is the figure given for these drugs, in people who took them as prescribed. The examination version of this error is the candidate who lists absence of dependence liability as an advantage of the class. The clinical version is the prescriber who discontinues a long-standing hypnotic without a plan, because the label said the drug was not a benzodiazepine.

24.7 The hypnotics that never touch the benzodiazepine site

The benzodiazepine receptor agonists are one family among several, and a comparison confined to them misses the axis that separates the classes. Practical Psychopharmacology sets out the wider field: beyond the benzodiazepine receptor agonists, the hypnotic families in use are the antihistamines, the **melatonin receptor agonists** and the **orexin receptor antagonists**. Read as a list of targets rather than a list of drugs, those are distinct ways of producing sleep.

- **GABA-A modulation.** The benzodiazepines and the Z-drugs, acting at the benzodiazepine recognition site, with the profile set out above.
- **Histamine blockade.** The antihistamines, several of which are available without prescription in many countries, producing sedation as an extension of a receptor action that is otherwise an unwanted effect.
- **Melatonin agonism.** Agents acting at melatonin receptors, working with the timing system rather than by depressing arousal.
- **Orexin antagonism.** Agents blocking orexin signalling, and therefore removing a wake-promoting drive rather than adding an inhibitory one.

That last distinction is the conceptual pay-off. A GABA-A modulator produces sleep by increasing inhibition; an orexin antagonist produces it by subtracting wakefulness; a melatonin agonist works through the circadian signal. Because the two newer families act nowhere near the benzodiazepine recognition site, the tolerance and withdrawal pharmacology set out above does not transfer to them by the same mechanism. What the sources quoted here do not supply are half-life or receptor-occupancy figures for the melatonin and orexin agents, so no kinetic comparison with the Z-drugs can honestly be printed, and none is attempted. The licensing position of individual agents also differs between countries, so the class names travel where product names do not. Their place in the treatment of insomnia belongs to the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes.

24.8 Comparing any two hypnotics

The comparison that examinations reward, and the one that survives the arrival of new molecules, runs in a fixed order. Establish the mechanism family first, because it decides whether the dependence and withdrawal pharmacology of the benzodiazepine site is in play at all. Within that family, establish the subunit preference, remembering that it is graded and that for the cyclopyrrolones it is disputed. Then take the half-life, which predicts how much drug is present when the patient must drive, work or supervise a child. Only then read the adverse-effect signature, which for this group includes psychomotor impairment that tolerance does not fully abolish and, for one agent, a characteristic taste disturbance.

MNEMONIC

Zaleplon, Zolpidem, Zopiclone, eszopiclone: shortest half-life first

The agents listed in ascending order of elimination half-life are also the order in which they are usually asked about, and that order predicts, other things being equal, how much drug is still present when the alarm goes off. **Za**leplon is the outlier at roughly an hour, **Zol**pidem is short but longer, and the cyclopyrrolones sit at the long end.

What none of this settles is whether a given patient should be taking a hypnotic at all. That decision belongs with the assessment of the sleep complaint. The pharmacology is what makes it executable: it tells the prescriber what the chosen molecule will still be doing at breakfast, what will happen when the patient stops it, and which of the reassuring things said about the class are properties of the drug rather than properties of its name.

The Z-drugs are benzodiazepine-site agonists whose prefix hides the fact: preference for the alpha-1 population is graded and strongest for zolpidem and zaleplon, half-life decides what is left in the morning, and withdrawal on stopping is common in people who took the drug exactly as prescribed.

25 · The non-benzodiazepine anxiolytics and the cognitive enhancers

Every drug gathered here is named for what it is not. The non-benzodiazepine anxiolytics are defined by a negative: they relieve anxiety without acting where a benzodiazepine acts, each by a different molecular route, so the group shares no receptor, no time course and no adverse-effect signature. The cognitive enhancers are the same kind of grouping seen from the other side, three drugs that inhibit an enzyme and one that plugs an ion channel, sharing a shelf rather than a mechanism. Treat either family as a single class and most questions about it become unanswerable; read each agent by its target and the answers fall out of the target.

The marks sit on dissociations, and three are worth naming in advance: a plasma half-life of a few hours against an onset measured in weeks; a plasma half-life that says nothing about how long an enzyme stays inhibited; and one mechanism the standard textbooks label with two different words. Whether any of these agents is the right treatment for a particular anxiety disorder belongs to the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes, and their

indication, staging and efficacy in dementia belong to the Dementias and neurocognitive disorders notes. Cross-class interaction tables and monitoring schedules belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

2 to 4 hours	10 days	1.6 hours	about 70 hours
BUSPIRONE PLASMA HALF-LIFE	BEFORE ITS ANXIOLYTIC EFFECT BECOMES APPARENT	RIVASTIGMINE MEAN ELIMINATION HALF-LIFE	DONEPEZIL HALF-LIFE, HENCE ONCE-DAILY DOSING

25.1 An azapirone, and the receptor it never touches

Buspirone is an **azapirone**, also written as an azaspirodecanedione, and the single most important fact about it is a negative one: it does not act on the GABA-A benzodiazepine chloride ionophore at all. That one statement generates most of what follows, because every property a benzodiazepine derives from modulating a chloride channel is one buspirone has no route to.

Where it does act is on a serotonin receptor. Buspirone is a **partial 5-HT1A agonist** and a relatively weak dopamine antagonist, apparently at predominantly presynaptic sites. The dopaminergic component is worth holding separately: it is the weaker action, and its site is qualified rather than asserted. What partial agonism does to a prescribing decision is developed where receptor theory is treated in this chapter, and the underlying receptor biology sits in the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

Two cautions belong with that receptor address. The Companion to Psychiatric Studies, having given the profile, describes buspirone's complex actions as incompletely understood, so the profile is a location and not a completed explanation. And candidates who have learnt that anxiolysis means a chloride channel will translate a serotonergic anxiolytic back into GABAergic language, then predict sedation, anticonvulsant activity and cross-tolerance that this drug does not have. The mechanism is not a paraphrase of the benzodiazepine mechanism; it is a different one that happens to earn the same therapeutic label.

25.2 A short half-life against a slow onset

The kinetics are unremarkable, and worth learning because they explain so little. Buspirone is rapidly absorbed, undergoes extensive presystemic metabolism and has a half-life of only **two to four hours**. A drug cleared that quickly rises and falls several times a day, which on the reasoning used everywhere else should predict a rapid, easily titrated effect.

It predicts nothing of the sort. Buspirone's action is not instantaneous, requiring some 10 days to become apparent and up to 3 weeks to develop fully. Set those two figures side by side and the arithmetic is unavoidable: the effect cannot be a simple function of plasma concentration, because the concentration has completed hundreds of cycles before the effect appears. Whatever produces anxiolysis here is downstream of receptor occupancy and is timed like an antidepressant response, not like a sedative.

The exact figure is not agreed. Tasman's psychiatry text puts the effect at least two weeks away and reports clinical experience suggesting three to four weeks before a truly beneficial effect is seen, while the Companion gives the shorter figure above. Nothing turns on which is right: the

teachable envelope runs from about ten days to about four weeks, and the contrast that earns the mark, against the very rapid onset of a benzodiazepine, survives either.

One vulnerability follows from the presystemic step. Buspirone's metabolism is inhibited by naringenin in grapefruit juice, which produces substantial increases in blood levels. That is a class-level property rather than an interaction table; the tables belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

TRAP **Quoting a single onset figure for buspirone.** Two standard texts give different numbers, and an answer key built on either will mark the other wrong. Say that the effect takes days to weeks rather than minutes, give the envelope, and name the contrast with the benzodiazepine; that answer is correct against both sources and is what the question is actually testing.

25.3 One effect where a benzodiazepine has four

Buspirone's spectrum of action is restricted to anxiety, with little or no hypnotic, anticonvulsant or muscle relaxant component. A benzodiazepine carries all four actions on one molecule because all four descend from the same channel; buspirone carries one, because it is not at that channel. The practical corollary is the hardest single rule here: it should not be used to substitute for benzodiazepines or alcohol in withdrawal regimes, and the same text warns that it has a mild proconvulsant action and should be avoided in those at risk of fits and in combination with monoamine oxidase inhibitors. Benzodiazepine dependence as a disorder belongs to the Substance use disorders notes.

That same absence of agonist activity puts buspirone at the far end of the misuse-liability spectrum. Buspirone is essentially devoid of benzodiazepine agonist activity and has little or no misuse liability in the predictive human models used to rank sedative-hypnotic drugs, placing it below the benzodiazepines and far below methaqualone and the barbiturates. Read the qualification with the claim: those rankings come from laboratory paradigms of drug liking and from the ARCI-MBG scale, not from field data, so what has been shown is an absence of reinforcing subjective effect, not an absence of misuse in the world.

Stopping the drug behaves accordingly. Discontinuation of buspirone has not been associated with the development of withdrawal symptoms, although patients discontinued from pharmacotherapy tend to relapse over time. That pairing is the honest version of the class advantage: relapse and withdrawal are different events, and a drug can be free of one while remaining fully subject to the other.

PITFALL

Starting buspirone to cover a benzodiazepine taper, then reading the patient's distress as anxiety returning. A drug with no action at the benzodiazepine site cannot occupy the receptor the patient is withdrawing from, so the taper runs unopposed while an agent that will not work for another two weeks is credited with holding it. The error is attractive because the two drugs share the word anxiolytic, and the proconvulsant caution makes it worse, since the population being tapered is the population with a lowered seizure threshold.

25.4 Beta-adrenergic blockade, the gabapentinoids, and onset as the axis

The other non-benzodiazepine routes are narrower than their reputation. Beta-adrenergic blockers such as propranolol may achieve short-term control of tremor and palpitations, which can be the most handicapping symptoms of specific social phobias such as performance anxiety among musicians, and it is doubtful whether they are more generally effective in social phobia. That sentence is the whole of what the Shorter Oxford Textbook of Psychiatry supports: two target symptoms, one narrow situation, and no mechanism at all.

Two things commonly taught alongside it are absent from the sources here, and both should be omitted rather than reconstructed. The familiar division of anxiety into somatic symptoms that a beta-blocker treats and psychic symptoms that it does not is not in this passage. Neither is the lipophilicity argument, the claim that propranolol is preferred over atenolol because it enters the central nervous system: no source here supports choosing a beta-blocker for its degree of central penetration, so that reasoning should not be offered as the pharmacological basis for the choice.

The gabapentinoids are the third route and the most often misfiled. Pregabalin binds to a subunit of voltage-gated calcium channels in the central nervous system, is unusually not protein bound and undergoes virtually no metabolism, so its interactions are few, and the Companion is explicit that how that binding relates to the drug's mode of action is unknown. The subunit is named elsewhere: **the alpha-2-delta protein** of the voltage-sensitive calcium channel is the target of pregabalin and gabapentin, and that protein may regulate conformational changes that alter how the channel opens and closes. The name gabapentinoid records a chemical resemblance, not a receptor: there is no action at the GABA-A receptor here, and a candidate who assumes one predicts the wrong interactions and the wrong withdrawal.

On a single axis of onset the group organises itself, and the axis is mechanism, not indication.

AGENT OR CLASS	MOLECULAR TARGET AS STATED	ONSET OF ANXIOLYTIC EFFECT
Buspirone	5-HT1A partial agonist	Delayed, at about two weeks
Beta-adrenergic blockers	Blockade of beta-adrenergic receptors	Rapid
Pregabalin	GABAergic by a mechanism not fully known	Rapid
Benzodiazepines	Modulation of the GABA-A receptor	Rapid to very rapid

25.5 The cognitive enhancers: two mechanisms, not one class

The second family is smaller than it looks and splits cleanly in two. Three cholinesterase inhibitors are currently available in the United States: donepezil, rivastigmine and galantamine, with tacrine, the first agent approved, no longer on that market. Standing apart from them, memantine is an NMDA receptor antagonist approved for moderate to severe Alzheimer disease. Both approval statements are specific to the United States and should be quoted as such. Which agent is indicated at which stage, and what either family achieves clinically, belongs to the Dementias and neurocognitive disorders notes.

The cholinergic three share a target that is more divisible than the class name suggests. Two enzymes hydrolyse acetylcholine, **acetylcholinesterase** and **butyrylcholinesterase**, and the cholinesterase inhibitors differ in the extent to which they selectively inhibit each. The second enzyme is the minor partner centrally, making up about 10 per cent of the total cholinesterase in the central nervous system according to the same source, which is why a drug that also blocks it is doing something additional rather than something more.

That split is what makes the within-class comparison tractable: there is no single quantity called cholinesterase inhibition to rank these drugs on, only two enzymes, a duration of binding, a nicotinic action that one agent has, and a set of kinetics.

25.6 What actually separates the three cholinesterase inhibitors

The enzyme selectivities are the first column and they are stated cleanly. Donepezil inhibits acetylcholinesterase with higher selectivity, rivastigmine inhibits both acetylcholinesterase and butyrylcholinesterase, and galantamine selectively inhibits acetylcholinesterase while also producing allosteric modulation of nicotinic receptors. The withdrawn agent completes the range in the opposite direction: tacrine inhibited butyrylcholinesterase 18 times more potently than acetylcholinesterase, which shows that selectivity within this class is a real variable.

Galantamine's second action is a genuinely different mechanism and is examined as such. **Allosteric modulation of nicotinic receptors** is distinct from the action of acetylcholine and from that of nicotinic agonists or antagonists, and it upregulates or potentiates nicotinic ion channel activity in response to acetylcholine, so cholinergic transmission is enhanced by a route that never passes through the enzyme.

The second axis is time on the enzyme rather than affinity for it. Cholinesterase inhibitors are classified as reversible, pseudo-irreversible or irreversible according to how long they bind the enzyme, and the three arms are worth separating explicitly.

- Reversible inhibitors bind the enzyme for seconds to minutes, so inhibition tracks the plasma concentration closely.
- **Pseudo-irreversible inhibitors** bind for hours, so inhibition outlasts the concentration that produced it.
- Irreversible inhibitors sit at the far end of the same axis, where the enzyme is not available again until new enzyme has been made.

MNEMONIC

Donepezil one, rivastigmine both, galantamine plus nicotinic

Three agents, three answers to one question about the target. **Donepezil** acts on acetylcholinesterase preferentially. **Rivastigmine** takes both enzymes. **Galantamine** is selective for the same enzyme as donepezil but adds a nicotinic action nobody else in the class has. The line records what each one does, never which is better.

25.7 Kinetics, and why the plasma half-life is the wrong question

The elimination half-lives inside this small class span a very wide range, and each must be read against the binding classification rather than on its own. Rivastigmine is rapidly absorbed orally with a mean time to maximum plasma concentration of one hour, range 0.8 to 1.7 hours, and has a mean elimination half-life of **1.6 hours**. Read naively, that predicts a drug present for only part of the day. It does not, because rivastigmine's inhibition is pseudo-irreversible, and for such an inhibitor the plasma half-life and the duration of enzyme inhibition are different quantities that need not resemble each other.

Galantamine sits in the middle and carries the class's one metabolic complication. Galantamine immediate release has an elimination half-life of about seven hours, and galantamine is mainly metabolised by hepatic CYP2D6 and CYP3A4, with the CYP2D6 metabolite sanguinine accounting for 20 per cent of the parent drug and being three times more potent than the parent drug at inhibiting acetylcholinesterase. An active metabolite more potent than its parent, formed by a polymorphic enzyme, is the classic setting for a genotype-driven difference in effect, though the source's own position is that the poor-metaboliser difference is probably not significant given the large variation between patients. The enzymology sits in the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes, and the cross-class interaction consequences in the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

The two long-acting agents anchor the other end. Donepezil's schedule is derived from its kinetics: the Indian pharmacology standard attributes its once-daily administration at bed time to a long half-life of about 70 hours, a distinct advantage over rivastigmine and galantamine, both of which need twice-daily dosing, and the Comprehensive Textbook of Psychiatry gives the same order of half-life for the same reason. Memantine is longer again: it reaches peak plasma concentration within three to seven hours, has an elimination half-life of **between 60 and 80 hours**, and is largely excreted unchanged in the urine, the remainder being converted to three minimally active polar metabolites, with steady state reached separately at about 21 days.

AGENT	ENZYME OR RECEPTOR ACTION	ELIMINATION HALF-LIFE	KINETIC CONSEQUENCE
Donepezil	Inhibits acetylcholinesterase with higher selectivity	About 70 hours	Once daily at bed time
Rivastigmine	Inhibits both acetylcholinesterase and butyrylcholinesterase	1.6 hours, mean	Twice-daily dosing
Galantamine	Selective for acetylcholinesterase, plus nicotinic allosteric modulation	About seven hours, immediate release	Twice-daily dosing
Memantine	NMDA receptor antagonist	60 to 80 hours	Steady state at about 21 days
Tacrine, withdrawn	Inhibited butyrylcholinesterase 18 times more potently than acetylcholinesterase	Not stated in these sources	Withdrawn from the United States market

GOOD TO REMEMBER

A half-life answers how often the drug must be given and how long it takes to reach steady state. It does not answer how long the target stays engaged, which for a pseudo-irreversible enzyme inhibitor is a separate measurement. With an agent whose steady state is three weeks away, anything assessed in the first fortnight measures the titration rather than the medicine.

RULE **A short half-life buys neither a fast onset nor a short action.** Buspirone is cleared in a few hours and takes weeks to work; rivastigmine is cleared faster still and inhibits its enzyme for longer than it is present. Onset is set by whatever process actually mediates the effect, and offset by how long the target stays engaged. The plasma curve sets the dosing interval and nothing else.

25.8 Memantine, and the block that works only while the channel is open

Memantine is the one agent here that is not a cholinergic drug, and its mechanism is the most precisely specified claim in the section. It is an **uncompetitive open-channel NMDA receptor antagonist**, and Stahl's account gives the four properties together: low to moderate affinity, voltage dependence, and fast blocking and unblocking kinetics, with the drug working at the same site as magnesium, plugging the ion channel where the magnesium ion normally blocks it at rest. Each clause does work. Because the drug occupies the channel rather than the glutamate recognition site, it can bind only once the channel has been opened, which is what voltage dependence means here; and because affinity is low and the kinetics fast, a barrage of glutamate from normal neurotransmission readily and quickly reverses the block.

The therapeutic logic follows directly. Chronic blockade of the NMDA receptor would hypothetically interfere with memory formation and neuroplasticity, so a drug meant to damp excessive tonic activation without abolishing phasic signalling has to be this kind of antagonist

and no other. Tasman's text describes the same mechanism independently, as an uncompetitive antagonist binding the NMDA receptor-operated cation channels and acting by inhibiting calcium influx.

One caution about the literature is worth carrying into the examination hall. At least one printing of a major psychiatry textbook carries a sentence that reads as though memantine blocked glutamate binding at the receptor. That is incompatible with open-channel block and with the surrounding paragraph, which describes selective uncompetitive blockade and strong voltage dependence. Where one sentence contradicts the mechanism its own paragraph describes, take the paragraph.

TRAP **Treating uncompetitive and noncompetitive as a factual disagreement about memantine.** Every source that spells the mechanism out in prose calls the drug uncompetitive, and the same books print noncompetitive in figure legends and passing mentions. The mechanism described is identical in both: binding inside the channel at the magnesium site, only while the channel is open, with low affinity and fast kinetics. Learn the mechanism, use uncompetitive, and treat a mismatch against an answer key on this one word as a labelling difference rather than an error of fact.

25.9 Nootropic: a claim about a category, not a mechanism

The word that ends this section is not a pharmacological description at all. Piracetam is a cyclic GABA derivative with no GABA-like activity, and the word **nootropic** denotes a drug said to selectively improve the efficiency of higher telencephalic integrative activities. Both halves repay attention. The molecule resembles a transmitter it does not act like, which is the trap the gabapentinoid name sets earlier; and the label names a hoped-for functional outcome rather than a target, so nothing about a receptor, an enzyme or a channel can be derived from it. A category defined by its intended effect is a claim awaiting evidence, not a mechanism.

The evidence has been examined and the verdict is unambiguous. Piracetam has long been promoted in India and some other countries for cognitive impairment and dementia in the elderly and for intellectual disability in children, but a Cochrane Database review of 2004 concluded that published data do not support such use, and in the United Kingdom it is approved only as adjunctive treatment of cortical myoclonus. The narrow surviving approval is the tell: the demonstrable effect was licensed, the broad cognitive claim that named the category was not.

IN INDIA

This is the one agent here where the prescribing decision is made against local practice rather than with it. Piracetam is promoted in this country for exactly the two populations in which the evidence was examined and found wanting, so it should not be started as a cognitive enhancer for an older adult with cognitive impairment or for a child with intellectual disability, and a request to add it to an existing regimen has no pharmacological argument behind it.

When a patient arrives already taking it, ask which claim it was started on. If the answer is cognition, say plainly that the licensed indication surviving review is a movement disorder rather than a cognitive one. The drug is well known in Indian households, and quietly stopping the tablet reads as a treatment being taken away.

Nothing in this section shares a mechanism with the drug it is offered as an alternative to, so nothing in it inherits that drug's speed, its spectrum or its behaviour on withdrawal.

26 · Psychostimulants and atomoxetine: onset, duration and diversion

What separates one stimulant preparation from another is almost entirely kinetic. The molecules in routine use are few: methylphenidate, dexamfetamine and the mixed amphetamine salts all raise synaptic catecholamines at the monoamine transporters. What differs between the products, and what is examined here, is rate. How fast the drug arrives, how long its effect lasts and how quickly it occupies a transporter determine both the shape of the therapeutic day and the probability that the tablet is swallowed by somebody other than the patient. Atomoxetine belongs in the same section for the opposite reason: it is not a stimulant at all, and set against the stimulants it demonstrates that abuse liability is a property of where and how fast a transporter is occupied rather than of which transmitter is raised. The mechanism itself, reuptake blockade against transporter-mediated release, and the indication-specific dosing, titration and adverse-effect monitoring of these agents, belong to the ADHD and specific learning/communication disorders notes; stimulant misuse as a disorder belongs to the Substance use disorders notes.

within 30 minutes	4 to 6 weeks	up to 12 hours	15 per cent
ORAL STIMULANT ONSET OF ACTION	ATOMOXETINE TIME TO EFFECT	LISDEXAMFETAMINE DURATION OF ACTION	METHYLPHENIDATE PLASMA PROTEIN BINDING

26.1 The one property the whole class is organised around

The class begins at the gut wall. Tasman Psychiatry puts it simply: psychostimulants are rapidly absorbed from the gut and therefore act quickly, an oral dose beginning to work **within the first 30 minutes** of ingestion. Almost nothing else in psychiatric pharmacology behaves like this: an antidepressant or an antipsychotic is judged over weeks, a stimulant before the patient has left the waiting room.

Three consequences organise the rest of the topic. First, the therapeutic trial is short, so response and non-response are both visible quickly rather than inferred from a plasma level. Second, because the effect appears and disappears within a single day, the clinically interesting number is

not the elimination half-life but the **effect duration**, the length of time a single administration continues to do useful work. Third, a drug that reaches the brain fast enough to be felt can by that same property reinforce its own consumption. Speed is at once the therapeutic advantage of the class and the reason it is regulated.

26.2 Onset, peak, half-life and duration are four different numbers

Residents lose marks by collapsing four independent parameters into one. Tasman Psychiatry tabulates them separately for the immediate-release preparations.

IMMEDIATE-RELEASE PREPARATION	EFFECT ONSET (MINUTES)	ELIMINATION HALF-LIFE (HOURS)	EFFECT DURATION (HOURS)
Methylphenidate	30 to 60	2 to 4	3 to 5
Dexamfetamine	30 to 60	3 to 5	4 to 6
Mixed amphetamine salts	30	3 to 5	4 to 6

The time to peak concentration is the fourth parameter and it separates the preparations least: the same table gives a time to peak of one to three hours for methylphenidate, for dexamphetamine and for the mixed salts alike. What does separate them is the pair at the right of the table. Methylphenidate has both the shortest elimination half-life, at **two to four hours**, and the shortest effect duration; the amphetamine preparations sit a step longer on both axes. That within-class ranking is the durable teaching point, and it is why immediate-release methylphenidate has needed more frequent administration than an amphetamine to cover the same waking day. Notice too that duration is not read off the half-life: for methylphenidate the tabulated effect outlasts the tabulated half-life, as you would expect of a drug whose action depends on transporter occupancy rather than plasma concentration. A half-life predicts how long a drug takes to leave the body, not how long the patient feels better, and here the second question is the clinical one.

TRAP **Quoting one methylphenidate half-life as though the standard texts agreed, and then treating that figure as the duration of action.** Two errors travel together. The first is that Tasman Psychiatry is internally inconsistent about the elimination half-life: its parameter table gives two to four hours, while the prose of the same chapter reports that methylphenidate peaks in plasma at 2 to 2.5 hours and falls to half the peak after 3 hours. This is one book disagreeing with itself, and because the prose point estimate sits inside the tabulated range, the range is the safer figure to carry into an answer. Declare the discrepancy rather than resolving it silently. The second error is the more consequential: candidates who have memorised a short half-life conclude that every methylphenidate product needs several administrations a day, which is false the moment a modified-release preparation is prescribed. Half-life belongs to the molecule; duration of action belongs to the product.

26.3 Where methylphenidate actually goes

Methylphenidate has an unusual disposition for a psychotropic, and its interaction profile falls straight out of it. The parent compound is hydrolysed at its ester group to yield **ritalinic acid**, an inactive product, and approximately 20 per cent is oxidised by the liver to secondary metabolites.

The dominant route is an esterase reaction rather than a cytochrome one, which is why this drug does not sit at the centre of the psychotropic substrate map that the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes set out.

The second peculiarity is how little of it needs to be there. Its low plasma protein binding, of about 15 per cent, makes it highly available to cross the blood-brain barrier, which is part of why relatively low plasma concentrations are effective. The comparison the source draws is striking: clinically effective oral methylphenidate produces plasma concentrations as low as 7 to 10 nanograms per millilitre, against dexamfetamine concentrations between 40 and 120 nanograms per millilitre in children with ADHD. Such a gap suggests a large **first-pass effect**, and the reported oral bioavailability has itself moved: once given in the 80 per cent range, it is placed by more recent studies close to 30 per cent. A drug largely destroyed before it reaches the circulation, which nonetheless works an order of magnitude below its comparator, is one whose plasma concentration tells you very little.

The site of that metabolism is contested inside a single book, and the honest position is to say so. Tasman Psychiatry states in prose that methylphenidate is primarily metabolised extrahepatically by **de-esterification** to an inactive metabolite, while the parameter table of the same chapter records its metabolism as hepatic, and the passage above assigns a fifth of the drug to hepatic oxidation. Teach the mechanism, not the compartment: clearance is mainly by de-esterification to inactive ritalinic acid with a minor oxidative route, and the extrahepatic claim is not settled enough to offer as though it were.

26.4 The release mechanism, not the molecule, sets the duration

Formulation is the principal therapeutic variable in this class, and it is engineering rather than chemistry. The same molecule can be built to act for a third of a working day or for the whole of one, and the design that achieves it is stated on the product, not deduced from the drug. Tasman Psychiatry maps the **release mechanism** of each preparation onto the duration it buys.

PREPARATION	MECHANISM OF DRUG RELEASE	DURATION OF ACTION (HOURS)
Mixed amphetamine salts	Single pulse	4 to 6
Mixed amphetamine salts, long acting	Two pulses	8 to 10
Methylphenidate, sustained release	Wax matrix	3 to 8
OROS methylphenidate	Three pulses	10 to 12
Transdermal methylphenidate patch	Transdermal	10 to 12
Lisdexamfetamine	Delayed single pulse	Up to 12

Read down the duration column. One pulse delivers a single bolus and buys the shortest cover; two pulses buy roughly double; three pulses, or a transdermal system that bypasses the gut and releases continuously, buy the longest. The wax matrix is the outlier: its band is the widest of any

design here, so such a product is the least predictable and the least suitable when the timing of cover matters. The prodrug reaches a long duration by another route, its delay being metabolic rather than physical.

PITFALL

Switching a patient between two preparations of the same drug on total daily dose alone. Methylphenidate is not one product but several, and the tabulated durations of a wax-matrix and an osmotic three-pulse preparation do not overlap at their extremes. A switch made on the drug name changes the shape of the day even when the total dose is unchanged, and the complaint that follows is usually loss of afternoon efficacy or a new evening problem, neither a failure of the molecule. Name the preparation, not just the drug, on every prescription and in every handover note.

26.5 Why speed, and not the molecule, is the abuse variable

The pharmacology of diversion is the core of this topic, and it is a story about rate. Tasman Psychiatry records that the slow reversal of methylphenidate binding to the dopamine transporter has been interpreted to mean that methylphenidate is not as reinforcing as cocaine, and is therefore less likely to lead to self-administration, even though therapeutic oral doses significantly increase extracellular dopamine. That last clause is the part candidates miss. The dopamine rises; reinforcement does not follow, because it tracks how abruptly occupancy changes rather than how much transmitter is present.

Stahl's Essential Psychopharmacology generalises the principle into a class statement, and flags it throughout as hypothetical. Extended-release oral stimulants, the transdermal methylphenidate patch and the prodrug lisdexamfetamine are considered **slow-dose stimulants**: three quite different pieces of engineering converging on one kinetic result.

- an oral extended-release system, which meters the drug out of the tablet;
- a transdermal patch, which meters it through the skin and never presents a gut bolus at all;
- a prodrug, which meters it by requiring an enzymatic step before any active drug exists.

The shared mechanism is regional. These preparations occupy noradrenaline transporters in the prefrontal cortex with slow enough onset, and for long enough duration, to enhance tonic signalling, but they do not occupy dopamine transporters in the nucleus accumbens quickly or extensively enough to increase **phasic signalling** at D2 receptors, and it is that failure to drive the accumbens phasically that hypothetically reduces abuse potential. Tonic cortical signalling is what the therapeutic effect is thought to need, phasic accumbens signalling is what reinforcement needs, and formulation supplies the first while withholding the second.

Lisdexamfetamine is the purest expression of the idea. It is a **prodrug** in which dexamfetamine is complexed with the amino acid lysine and is inactive in that form, and it is broken down in red blood cells so that dexamfetamine is gradually made available, which is why it is unlikely to be abused for recreational or dependency-driven purposes. Because the rate-limiting step lives inside the erythrocyte rather than the capsule, crushing, snorting or injecting the contents

defeats nothing: the route changes, the conversion step does not. The source applies the same practical logic to extended-release methylphenidate preparations.

RULE **Abuse liability in this class is set by the rate and the site of transporter occupancy, not by the identity of the drug or by the amount of dopamine released.** Every fact in this section applies that one sentence. Therapeutic oral methylphenidate raises extracellular dopamine and is still poorly reinforcing, because its binding to the transporter reverses slowly. A modified-release and an immediate-release preparation of the same molecule carry different risk, differing in nothing but rate. A prodrug is safer than its own active metabolite, because a conversion step is inserted between the mouth and the accumbens.

MNEMONIC

Slow in, long on, cortex not accumbens

Slow in: the onset must be gradual, which is what every extended-release, transdermal and prodrug design is engineered to achieve. **Long on**: occupancy must be sustained rather than spiky, because tonic signalling is what the therapeutic effect requires. **Cortex not accumbens**: the transporters to occupy are the noradrenaline transporters of the prefrontal cortex, not the dopamine transporters of the nucleus accumbens.

26.6 The regulatory expression of a kinetic fact

Regulation follows the pharmacology, and the examinable content is the reasoning rather than the schedule. The Maudsley Prescribing Guidelines record that both methylphenidate and dexamfetamine are Schedule 2 Controlled Drugs in the United Kingdom, that prescriptions should be written appropriately, with the total amount in words and figures, and that they should cover **a maximum supply of 30 days**. The same source makes a within-class discrimination worth carrying: dexamfetamine is probably more likely than methylphenidate to be diverted and misused. The source states the ranking without giving a mechanism, so carry it as an empirical comparison rather than as something deduced from the kinetics above, and comparative rather than a licence to regard methylphenidate as safe.

The scheduling quoted here is one jurisdiction's and should be labelled as such whenever it is used. The corresponding Indian position, the schedule these drugs occupy in Indian law and the prescription requirements that follow, is not established in the sources behind this section and is deliberately not asserted; take it from current Indian statute and current drug-controller requirements, which change more often than pharmacology does. What transfers across jurisdictions is the underlying claim: a preparation is regulated because of how fast it reaches the brain, so a preparation engineered to reach the brain slowly is making a regulatory argument as well as a therapeutic one.

26.7 Atomoxetine: the negative control for the whole argument

Atomoxetine is best taught as the experiment that isolates the variable. KDT Essentials of Medical Pharmacology defines it precisely: a **selective noradrenaline reuptake inhibitor**, unrelated to amphetamine and to imipramine, which does not enhance dopamine release in the brain, and is neither a central nervous system stimulant nor an antidepressant. Every element of

that definition works. It is not an amphetamine derivative, so structural analogy predicts nothing. It is not a tricyclic, so that class's receptor promiscuity does not apply. And it is not an antidepressant despite inhibiting a monoamine transporter, a useful reminder that mechanism does not confer class membership.

How it avoids being reinforcing is a matter of regional anatomy. Stahl's Essential Psychopharmacology explains that by blocking noradrenaline transporters, atomoxetine raises both noradrenaline and dopamine in the prefrontal cortex, where the inactivation of both transmitters is largely by those transporters, while the relative lack of noradrenaline transporters in the nucleus accumbens prevents it from raising either transmitter there, and that regional difference reduces the risk of abuse. This is the slow-dose argument by another road: the stimulant preparations avoid the accumbens by being slow, atomoxetine because the transporter it blocks is scarce there. Other noradrenaline transporter inhibitors would be expected to have the same effects.

Its disposition introduces the topic's one pharmacogenetic variable. Atomoxetine is absorbed orally, hydroxylated by CYP2D6 and excreted in the urine mainly as the glucuronide, and while the majority of individuals are extensive metabolisers, a few are **poor metabolisers** because of polymorphism of CYP2D6. The enzymology belongs to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes, and co-prescription with a CYP2D6 inhibitor to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes. What belongs here is the class-level consequence: a drug cleared through a polymorphic enzyme has a wider spread of exposure across a population than one cleared by an esterase.

The clinical headline is the timescale. The Maudsley Prescribing Guidelines give the time to effect for atomoxetine as **approximately four to six weeks**, and note in the same breath that it is not a controlled drug and may be useful where the diversion of stimulants is a problem. Held against an oral stimulant working inside the first half hour, that is an enormous gap in time to effect between two drugs given for the same purpose, and it is the comparison most likely to be examined here.

TRAP **Treating atomoxetine as a slow stimulant, and abandoning it early because nothing has happened.** The stem describes a patient started on atomoxetine who is unchanged after a fortnight, and offers a switch or a dose increase as the tempting answer. Both are premature. A stimulant declares itself the same day, so an unhelpful stimulant trial is over quickly; a noradrenaline reuptake inhibitor is assessed over weeks, and the quoted interval has not elapsed. The deeper version is classificatory: candidates who file atomoxetine under stimulants expect it to carry stimulant liability, and are then unable to explain why it is the agent recommended where diversion is the problem.

GOOD TO REMEMBER

Three questions convert this section into a prescribing conversation. What does the day need to look like, which is a question about effect duration and therefore about the release mechanism rather than the drug name. What happens to the tablets between the clinic and the patient, which kinetics answers: a prodrug or a patch where supervised immediate-release dosing cannot be guaranteed. And how long is everybody prepared to wait, because a family told that timescale in advance behaves very differently from one told nothing.

In this class the kinetics are the pharmacology: onset, effect duration and half-life are separate numbers that must never be substituted for one another, the release mechanism rather than the molecule sets how long a preparation works, and both the therapeutic effect and the abuse liability follow from how fast and where transporter occupancy changes, which is why a prodrug, a patch and a non-stimulant reuptake inhibitor all reduce diversion risk by the same argument.

Consolidate and apply

27 · Worked vignettes

A pharmacology vignette asks what the drug is doing, not what the diagnosis is. The examiner rarely asks a candidate to define a receptor. Far more often a short stem describes a patient already on treatment, or about to start it, plants one pharmacological fact, and waits to see whether the candidate reasons from it. That planted fact is almost always one of four things: a binding profile, an elimination half-life, an enzyme that some other drug has just inhibited or induced, or a margin between the concentration that works and the concentration that harms. This closing section teaches no new agent and no new class. It rehearses the pharmacology assembled across this chapter as a reading skill, moving from a paragraph of history to the mechanism that explains it.

The discipline that earns marks here is resisting the reflex to answer with a treatment algorithm. A stem that describes a patient whose sedation is intolerable is not asking which antipsychotic treats schizophrenia; it is asking which receptor family the current agent occupies and what a within-class move would change. Where a stem genuinely asks which drug to choose for a disorder, the reasoning belongs to that disorder's own chapter, and the antipsychotic and antidepressant selection decisions are set out in the Schizophrenia: management, antipsychotics and adverse effects notes and the Mood disorders: pharmacotherapy notes respectively. What this chapter supplies, and what a good vignette answer uses, is the layer underneath those decisions.

27.1 Four questions read almost any pharmacology stem

Before deciding what a stem is asking, run the same four questions across it. They are the four axes on which psychotropic agents actually differ, and between them they account for most of the

discriminating information an examiner can plant in a short paragraph.

Binding	Duration	Company	Margin
WHAT IT OCCUPIES	HOW LONG IT LASTS	WHAT ELSE IS ON BOARD	HOW MUCH ERROR IT FORGIVES

Binding asks which receptors and transporters the agent touches and with what relative affinity, because a drug's unwanted effects are largely the affinities that were not the point of prescribing it. Duration asks how long the drug persists, which governs how quickly it accumulates, how forgiving it is of a missed dose, and how quickly anything at all can be judged. Company asks what else the patient is taking, including the drugs a psychiatrist did not prescribe and the habits that behave like drugs, because a concentration can move without anyone changing the prescription. Margin asks how much distance sits between a useful concentration and a harmful one, which decides whether the drug can be managed by clinical judgement alone or needs a measured level. The single feature a stem plants to make one of these questions decisive is its **pharmacological discriminator**, and well-written stems are built around exactly one.

RULE **In a pharmacology stem, the answer is the mechanism, not the diagnosis.** If your answer to a stem in this territory can be written without naming a receptor, a half-life, an enzyme or a therapeutic margin, you have almost certainly answered the wrong question. The diagnosis in the stem is usually scenery: it tells you which class the patient is on, and nothing more. What earns the marks is the pharmacological property that makes this patient's problem inevitable and the within-class or within-schedule move that follows from it.

27.2 The vignette that resolves on a binding profile

The commonest stem in this territory describes a patient who is responding to treatment but cannot tolerate it. She is sedated all morning, or has gained weight steadily, or feels her heart racing when she stands, or has a dry mouth and constipation, or has developed galactorrhoea, or cannot sit still. Every one of those complaints is a receptor being occupied, and the reading move is **reverse inference**: take the symptom the patient reports, ask which receptor family generates it, then ask whether the prescribed agent has meaningful affinity there. The full map from binding profile to predicted adverse effect is developed in its own right within this chapter, and the biochemistry of the transmitter systems underneath it belongs to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

What makes the stem answerable rather than merely descriptive is that the receptor families cluster into recognisable symptom groups. Histaminergic and cholinergic occupancy produce one kind of complaint, adrenergic occupancy another, dopaminergic occupancy a third, and the serotonergic actions a fourth. Once the symptom has been mapped back to a family, the answer is a within-class move to an agent whose affinity at that family is lower, not a change of class and not a change of diagnosis. That is why the comparative profiles assembled elsewhere in this chapter are worth learning as a ranked shape rather than as a list of drugs: the vignette is asking you to substitute along one axis while holding the others.

TRAP **Treating members of a class as interchangeable because they share a label.** Candidates who have learned the classes as boxes will answer that any agent in the box will do, and lose the mark. Agents within a single class differ in what else they bind, how long they last and which enzymes clear them, and those differences are precisely what the stem is testing. **Binding promiscuity** is the reason two drugs with the same primary mechanism produce entirely different complaints, and the reason a within-class switch is a real therapeutic manoeuvre rather than a cosmetic one.

27.3 The vignette that resolves on a half-life

A second family of stems turns entirely on how long the drug stays. The patient misses a few doses and becomes unwell in a characteristic way; or stops the drug abruptly and develops symptoms that could be relapse or could be discontinuation; or has a level checked soon after a dose change and the clinician acts on it; or is switched from one agent to another and something goes wrong in the overlap. In every case the discriminator is the elimination half-life and the accumulation that follows from it, developed in full among this chapter's pharmacokinetic principles, with the underlying kinetic derivations in the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

Three consequences do most of the work in vignettes. A drug that is eliminated slowly buffers itself: missed doses matter less, the plasma concentration changes gradually, and symptoms on stopping are muted and delayed, so a picture arising soon after cessation is more likely to be something other than withdrawal. A drug eliminated quickly does the opposite, producing interdose symptoms, an abrupt **offset** and a withdrawal picture that arrives fast. And accumulation takes several half-lives whatever the drug, which means any judgement made before the concentration has plateaued, whether clinical or laboratory, is a judgement about a moving number. The interval a clinician must leave empty between stopping one agent and starting another, the **washout**, is the same arithmetic used in the opposite direction, and it is the reason some switches are simple and others are not.

PITFALL

Sampling a drug level before the concentration has plateaued, and then changing the dose on the result. A level drawn too early after starting or after a dose change is not a steady-state level, it is a point on a rising curve, and treating it as though it were the patient's true exposure leads to a dose increase in a patient who was already on their way to an adequate concentration. The same error in reverse explains apparently unstable levels in a patient whose adherence is intermittent. The remedy is arithmetic, not repetition: know the drug's half-life, know when the plateau is due, and sample against it.

27.4 The vignette that resolves on an enzyme

The third family of stems is the one where nothing about the psychotropic prescription changed and the patient changed anyway. A physician added an antimicrobial or an antifungal; a cardiologist started something new; the patient bought a herbal preparation over the counter; the patient stopped smoking on admission to a ward where smoking is not allowed. Then a drug that had been tolerated for months becomes toxic, or a drug that had been working stops working, and the stem waits to see whether the candidate looks past the psychiatric drug list.

The direction of reasoning is simple and it is worth having automatic. An added inhibitor reduces the clearance of a substrate, so the substrate's concentration rises, and it rises promptly, because inhibition acts at an enzyme that is already present. An added inducer increases the amount of enzyme, so the substrate's concentration falls, and it falls with a lag, because new enzyme has to be synthesised, and it recovers over a comparable but somewhat longer interval when the inducer is withdrawn, the figures for both directions being carried by this chapter's cytochrome P450 section. That asymmetry of tempo is itself a discriminator: an abrupt deterioration points toward inhibition, a loss of effect that develops over days rather than hours toward induction. The clinically important variant is **phenoconversion**, in which a patient with entirely normal metabolic capacity is turned by a co-prescribed inhibitor into the functional equivalent of a poor metaboliser. The psychotropic substrate map itself is developed within this chapter; the interaction matrix spanning several classes at once belongs to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

GOOD TO REMEMBER

When a stable patient destabilises, date the change against the whole drug list, not the psychiatric one. Ask what a non-psychiatric colleague started or stopped, what the patient bought without a prescription, and whether tobacco use changed, because admission to hospital silently removes an inducer from many patients. A symptom whose onset lines up with a change somewhere else in the list is an interaction until proved otherwise, and the question that resolves it is which enzyme both drugs share.

27.5 The vignette that resolves on the margin

Some stems are constructed so that the pharmacology is unremarkable except for one property: the drug has very little distance between the concentration that treats and the concentration that harms. Then a modest kinetic disturbance, one that would be trivial for a forgiving drug, carries the patient across the line. A **narrow therapeutic index** is what converts an ordinary intercurrent illness, a change in fluid balance, or a new co-prescription into a toxic presentation, and it is the single property that justifies routinely measuring a concentration at all. This chapter develops the index itself, and which psychotropics need a measured level, as a principle; the target concentrations, and the monitoring framework built on them, belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

The reading move is to notice that the stem has told you about the patient's physiology rather than about the psychiatry. Anything that alters volume status, reduces renal clearance or competes for the elimination route will move the concentration of a narrow-index drug while the prescription stays identical, and for a drug with a wide margin the same events pass unnoticed. Which precipitants matter for which agent is developed within this chapter's sections on the narrow-index drugs, and the answer a stem wants is rarely to stop the drug reflexively; it is to name the kinetic precipitant, state that the concentration will have moved, and act on that. Toxic presentations themselves, and the psychotropic emergencies more broadly, are recognised and managed under the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

27.6 The vignette that resolves on intrinsic activity

A subtler stem describes a patient in whom a drug was added and something got worse rather than better, or in whom two agents acting on the same system produced an effect neither would have produced alone. The property being tested is **intrinsic activity**: whether the agent at a given receptor is a full agonist, a partial agonist, a plain antagonist or an inverse agonist. A partial agonist behaves in opposite directions depending on what the receptor is already seeing, damping the system where tone is high and supporting it where tone is low, which is why adding one to a patient already maintained on a full antagonist can produce an apparent loss of effect. This chapter teaches partial agonism as a prescribing variable in full, and the receptor theory beneath it belongs to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

The same reasoning explains a second stem shape, in which an agent modulates the action of an endogenous ligand rather than blocking it, so its behaviour depends on the state of the system rather than on the dose alone. With this axis in hand, the question is not how much drug is present but what the receptor was doing before it arrived.

27.7 The vignette that resolves on tempo

When a problem starts is often more informative than what it looks like. Adverse effects have characteristic latencies, and the **latency** a stem specifies is frequently the discriminator that assigns the phenomenon to a mechanism. Something that appears within hours of a first exposure, before any accumulation has occurred, is a direct pharmacodynamic effect of occupancy. Something that emerges over the first weeks, as the concentration plateaus and adaptive changes begin, is typically dose-related and often modifiable by dose. Something that appears only after months or years of continuous exposure implies receptor adaptation rather than occupancy alone, and is the tempo of the late movement disorders whose spectrum this chapter develops in full. Something whose onset bears no relation to dose or duration at all is **idiosyncratic**, and cannot be titrated around.

Tempo also separates the two commonest look-alikes in this whole territory. Symptoms that appear within a short and predictable interval of stopping a drug, an interval read against that drug's half-life, and that resolve when the drug is reinstated, are a withdrawal or discontinuation phenomenon. Symptoms that appear later, at a remove the kinetics do not account for, and that do not remit promptly on reinstatement, are relapse. Do not separate them by what they look like. Where a drug was prescribed for the very state it suppresses, withdrawal reproduces the original presentation almost exactly, a trap the benzodiazepine section works through in detail. The discrimination is kinetic and pharmacological: the interval from cessation, and the response to reinstatement. Where the phenomenon is a substance use disorder in its own right rather than a pharmacological adaptation, the assessment and treatment belong to the Substance use disorders notes, and the clinical use of the anxiolytics and hypnotics belongs to the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes.

27.8 Look-alike stems and the reading spine

Consolidation is easiest as pairs. The table sets out the stems most often confused, against the property that separates them and the axis the reasoning lands on.

TWO STEMS THAT LOOK ALIKE	THE PHARMACOLOGICAL DISCRIMINATOR	THE AXIS IT LANDS ON
Symptoms soon after stopping a drug against symptoms emerging later, the two pictures looking much alike	Interval from cessation, and whether reinstatement resolves it	Half-life and offset
Loss of effect over weeks against abrupt toxicity in a stable patient	Whether the added agent induces or inhibits	Enzyme interaction
Intolerable sedation against intolerable restlessness on the same class	Which receptor family the agent occupies most avidly	Binding profile
An apparently unstable drug level against a genuinely rising one	Whether the sample was drawn at plateau	Accumulation kinetics
Toxicity after an intercurrent illness against toxicity after a dose increase	Whether a kinetic precipitant altered clearance	Therapeutic margin
Worsening after adding a second agent to an established one against worsening on the first alone	Full antagonism against partial agonism at the shared receptor	Intrinsic activity
An adverse effect present from the first exposure against one emerging after prolonged treatment	Latency, and its relation to dose	Tempo of onset

MNEMONIC

What it binds, How long it lasts, Alongside what else, Tolerance of error

The four questions to run across any pharmacology stem: **W**hat does the agent bind, and with what relative affinity; **H**ow long does it last, and has it reached plateau; **A**longside what else is the patient taking, prescribed or not; and what **T**olerance of error does the drug allow between a useful and a harmful concentration. They spell WHAT, which is the question the stem is really asking, and each one is a door to a different section of this chapter.

With the discriminator identified, a marks-winning answer follows four moves in order.

- Name the pharmacological property the stem has planted, in pharmacological language.
- State the mechanism that connects that property to the patient's problem.
- Choose the move that acts on that mechanism: a within-class substitution, a change of interval or schedule, a washout, or a measured concentration.

- Say explicitly what you would not do, and why the obvious algorithmic answer does not address the mechanism.

That final move is where the pharmacology becomes visible to an examiner. A candidate who says the drug should be stopped has answered a question about a patient; one who says it should not, because the problem is an inhibitor that will wash out or a withdrawal effect that reinstatement will resolve, has answered a question about pharmacology. Every vignette here rewards the second answer, worked from the four axes without a treatment algorithm.

> In a pharmacology vignette, identify the property the stem has planted - binding, duration, company or margin - and answer with the mechanism it implies, not with a treatment algorithm.

28 · Consolidation and quick review

Every class in this chapter holds several agents that share a mechanism and differ in what they cost the patient. That sentence is the chapter compressed. Each class was assembled around one shared action: blockade of the dopamine receptor, inhibition of monoamine reuptake, potentiation of the inhibitory chloride channel, inhibition of the metabolising enzyme. The shared action makes the agents members of a class and explains why they work at all. Everything else about them, what else the molecule binds, how long it stays in the body, and how much room lies between the exposure that helps and the exposure that harms, is where the members separate. A reader who knows the class knows why the drugs work and still cannot say which of them to give.

This closing section teaches no new agent, because every drug has already been taught in the section that owns it. What remains is to state the argument those sections were arranged to make, and to convert it into something that can be produced from a blank page. The argument is that **within-class selection** is a prediction rather than a preference: given a drug's binding profile, its handling by the body and the width of its safety margin, most of what separates it from its neighbours in the same class can be worked out before any of it is observed in a patient. The pharmacology needed to make that prediction is small, finite, and the same for every class here.

28.1 One argument, tested five times

The chapter moves through receptor and kinetic principles, then the antipsychotics, the antidepressants, the mood stabilisers, and the anxiolytics, hypnotics and remaining agents. Read as five territories they are a long list to be survived. Read as one argument they are a single claim tested five times, and the claim is that the difference between two members of a class is generated by pharmacology that can be stated in advance. The principles are not an introduction to the four drug bands; they are the instrument those bands are read with.

The claim has a testable shape, which is what makes it worth holding. Take two agents from one class and set out, for each, what it binds beyond the shared target, how quickly it is cleared and by what route, and how much distance lies between the dose that works and the dose that harms.

Those three statements anticipate most of what the patient reports, most of what the prescriber has to manage, and most of the practical difference in how the two drugs are started, stopped and exchanged. What they do not anticipate is which agent a unit happens to stock, or which one the reader saw used most often as a junior, and those two considerations decide a great many real prescriptions.

Each pairwise difference runs along a **separating axis**: one pharmacological property on which the members of a class can be ranked against each other. The axes are few and they recur. Sedation, weight, prolactin and movement liability rank the antipsychotics. Half-life, enzyme inhibition and binding beyond the transporter rank the antidepressants. Index width, elimination route and reproductive risk rank the mood stabilisers. Receptor subunit selectivity and half-life rank the hypnotics. Naming the axis rather than the drug converts a remembered preference into a defensible choice.

RULE **The shared mechanism explains the effect; everything that separates the members of a class is a side action, a kinetic difference or a margin.** Establish what two agents have in common only to confirm they belong to the same class. The clinical question is always what they do not share, and that question has a small and recurring set of answers.

28.2 Five questions asked of every agent

Because the separating axes recur, a single **retrieval card** serves the whole chapter. Rather than storing each agent as an unstructured block of facts, the reader learns one set of prompts and answers them for the drug in hand, which reduces the memory load and guarantees that no limb of an answer is dropped under time pressure. They spell CLASS, which is the level this chapter works at and the level at which an answer here should be pitched.

MNEMONIC

Clearance, Ligand behaviour, Adverse-effect signature, Safety margin, Selection

For any agent in any class, retrieve: **C**learance, meaning the half-life, the route of elimination and the enzymes that handle the drug; **L**igand behaviour, meaning whether the molecule acts as a full agonist, an antagonist, a partial agonist, a reuptake inhibitor or an enzyme inhibitor, and whether that action is reversible; **A**dverse-effect signature, meaning what else the molecule binds and therefore what the patient will report; **S**afety margin, meaning the distance between the exposure that treats and the exposure that harms; and **S**election, meaning the within-class choice that all four preceding answers were gathered to justify. Read together they spell CLASS, which is exactly the level this chapter teaches.

The card is built out of constructs the chapter has already supplied, so nothing in it has to be learned separately. Its value is that it disciplines an answer. A reader who has produced all five prompts for the agent in front of them has generated a structured response: how the drug is handled, what it does at the receptor, what the patient will feel, how much room for error exists,

and what follows for the choice. An examiner reads that as pharmacological reasoning rather than as recall of a list.

Binds	**Stays**	**Margin**	**Chooses**
WHAT ELSE IT TOUCHES	HOW LONG IT LASTS	ROOM BEFORE HARM	THE AGENT THAT FOLLOWS

28.3 What each question actually predicts

The card is only useful if each prompt is attached to the decision it settles, because a prompt that returns a fact without a consequence is one more thing to memorise. Set out this way, the four pharmacological questions each buy a different practical answer and the fifth is the answer itself. The table is a retrieval index rather than a teaching table: the full account of each property lives in the section that owns it.

THE PROMPT	WHAT THE ANSWER PREDICTS	WHERE MEMBERS OF A CLASS DIVERGE
Clearance: half-life, elimination route, metabolising enzyme	Dosing interval, time to a steady concentration, whether a measured level means anything yet, how abruptly the drug may be stopped, and how a switch is staged	Agents sharing a mechanism still differ on stopping and on switching, because the shorter-acting member falls away faster
Ligand behaviour: agonist, antagonist, partial agonist, reuptake inhibitor, enzyme inhibitor, and reversibility	Whether the effect has a ceiling, how the system behaves on withdrawal, and whether one molecule can look stimulating in one state and blocking in another	A partial agonist and a full antagonist at the same receptor share a class and are not interchangeable
Adverse-effect signature: what else the molecule binds	The symptoms the patient will report, anticipated before any appear, and the tolerability conversation held before the first prescription	Promiscuous binding predicts a broad adverse-effect signature; a cleaner profile predicts a narrower one at the same efficacy
Safety margin: distance between effective and toxic exposure	Whether the drug is titrated to clinical effect or to a measured concentration, and how much a small change in fluid balance or renal handling matters	A narrow margin makes an ordinary event dangerous where a wide margin absorbs it
Selection: the member that fits this patient	The actual prescription, defended by the four answers above rather than by habit	This is the output, not another property to memorise

Worked in both directions, given an agent produce its properties and given a property produce the agents at either end of it, the index becomes a network rather than a sequence.

28.4 Retrieve across the classes, not down one

A chapter learned in the order it was read is fragile, because recall then depends on starting at the top of a class and running down it. Durable retrieval comes from asking the same question of every class in turn, which is **interleaved retrieval**: deliberately mixing the classes during practice

instead of drilling one to exhaustion. It feels slower than working through a single class at a time, and that difficulty is why it works, because it trains the discrimination a stem actually tests.

Two prompts are worth interleaving across all five sub-areas. The first is the separating axis. The second is the **class-defining hazard**: the adverse consequence characteristic of the whole class rather than of one member, which therefore has to be considered whenever any member is prescribed. Together they explain both why a class is used and why choosing carelessly inside it has consequences.

CLASS	WHAT SEPARATES THE MEMBERS	THE CLASS-DEFINING HAZARD
Antipsychotics	Sedation, weight, prolactin and movement liability, each read off the binding profile rather than off the older typical and atypical labels	Movement disorder from dopamine blockade, and metabolic burden from the wider binding profile
Antidepressants	Half-life, inhibition of the metabolising enzymes, and how much the molecule binds beyond its intended transporter or enzyme	Disturbance on stopping, which tracks the half-life, and interaction risk, which tracks enzyme inhibition
Mood stabilisers	Width of the therapeutic index, route of elimination, whether the agent induces its own metabolism, and reproductive risk	Toxicity in a narrow-index agent, and harm to a developing fetus with the agents that carry it
Anxiolytics and hypnotics	Receptor subunit selectivity and half-life, which together set both the daytime residue and the rebound	Tolerance, dependence and withdrawal as pharmacological consequences of sustained receptor exposure
Stimulants and the non-stimulant agents	Speed of onset, duration of action, and what the formulation does to the concentration curve	Misuse and diversion, which follow rapidity of onset rather than the molecule alone

As a drill this means posing the class as the question and producing the axis and the hazard from memory, then reversing it: given an axis, name the class it ranks; given a hazard, name the receptor action that generates it. That movement, from a mechanism to its consequences and back, is what a viva is built to elicit.

28.5 The failure mode this chapter exists to correct

Set against the pharmacological method is the alternative most prescribing actually runs on, the **familiarity heuristic**: choosing the agent one has used before, seen used, or been taught by whoever supervised the first ward. Familiarity is not worthless, since a drug the prescriber knows well is monitored better than one they do not. It is simply not a reason, and it cannot be defended when the choice turns out badly. The remedy is not to abandon experience but to state the axis experience was tracking.

The examination version of the same failure is answering a within-class question with a name instead of a property. The stem describes a patient who cannot tolerate sedation, or who is

unlikely to attend for a blood test, or in whom a missed dose is likely; the answer that scores states the axis first and reaches the agent second. The antidote is **cold retrieval**, producing the card from a blank page with the book closed rather than reviewing an answer already written, because recognising a correct choice on the page is a far weaker skill than generating one from a bare stem.

TRAP **Answering a within-class question with a drug name instead of the property that selects it.** Candidates do this because the name is what they revised and because it feels decisive. The examiner is testing whether the choice can be derived, so an answer that names the separating axis and then lands on an agent earns marks an unsupported name does not, even when both arrive at the same drug.

PITFALL

Switching within a class and expecting the troublesome adverse effect to disappear, when the two agents share the binding that generated it. A switch is worth making only when the profiles genuinely differ on the axis causing the problem; where they do not, the patient endures a second titration and arrives at the same complaint. Before any within-class switch, state which property is expected to change and by what mechanism. If that sentence cannot be completed, the switch is being made on familiarity.

28.6 Where this chapter stops

Consolidating a chapter includes consolidating its edges, because a disciplined answer is defined as much by what it hands over as by what it teaches. The **handover boundary** matters unusually here, since the systematic pharmacology of a class sits beside several chapters that treat the same drugs from a different angle. Each of the following should be retrieved as an interface: named where it touches the question, handed to its owner, and left there.

- The underlying science of receptor theory, the kinetic equations and the enzymology of drug metabolism belongs to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes; this chapter uses those principles as instruments rather than deriving them.
- Choosing and managing an antipsychotic for a patient with a psychotic illness, including the management of movement disorder in that setting, belongs to the Schizophrenia: management, antipsychotics and adverse effects notes.
- Choosing an antidepressant or a mood stabiliser for a depressive episode or for bipolar disorder, and the clinical monitoring that goes with it, belongs to the Mood disorders: pharmacotherapy notes.
- The clinical use of anxiolytics and hypnotics belongs to the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes, while sedative dependence as a diagnosable disorder, with its assessment and detoxification, belongs to the Substance use disorders notes.
- Indication-specific dosing for attention deficit hyperactivity disorder belongs to the ADHD and specific learning/communication disorders notes; the antiepileptic lens on the

anticonvulsant mood stabilisers belongs to the Epilepsy and psychiatry notes; and the cognitive enhancers as treatments belong to the Dementias and neurocognitive disorders notes.

- Interactions spanning classes, prescribing in physical illness and in pregnancy, monitoring schedules running across psychotropics, therapeutic target ranges, and the drug emergencies all belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes.

Retrieved as a rehearsed set, the boundary supplies a ready-made handover sentence: name the neighbouring construct, say why it matters to the case in front of you, point to where it is taught, and return to the pharmacological question actually asked.

28.7 Weighting the comparison to the practice you will work in

Not every within-class comparison carries the same weight at the bedside you will actually stand at. Consolidation is finite, so the last allocation decision is where to spend the retrieval that remains. Weight it toward the comparisons that recur locally and the agents genuinely available to prescribe, rehearsing those pairs until the axis separating them comes without effort, and let rarely used agents sit at lower priority.

IN INDIA

Two pharmacogenetic differences taught in the sections that own them change a within-class decision on Indian ground, and both should be retrieved as decisions rather than as facts. Where a variant that slows clearance is documented as common in Asian populations, the decision is to open at the lower end of the range and titrate against response rather than to start at a standard dose and react to the adverse effect afterwards. Where a genetic test is recommended before a specific anticonvulsant is started in people of South-East Asian ancestry, the decision is to test before prescribing rather than after a reaction has occurred. Both remain within-class decisions, because in each case a differently handled member of the same class stays available.

28.8 Carrying the frame into the examination

If cold retrieval is the engine, **spaced retrieval** is the schedule that makes it last: returning to a class after a deliberate gap rather than in one continuous block. A memory retrieved just as it begins to fade is strengthened more than one retrieved while still fresh, so letting a class cool between attempts is a feature of the method rather than neglect. Across a revision period the gaps before returning to a secure class lengthen while a shaky one is brought forward, which concentrates limited hours on the gaps the reader actually has.

GOOD TO REMEMBER

Finish every within-class decision with a spoken sentence: this agent rather than that one, because of this property, at this cost. Said aloud in clinic it exposes a choice made on familiarity in a way that thinking it silently does not, and the same sentence is the one that earns marks in a viva. A decision that cannot be finished in that form has not yet been made pharmacologically.

Carried into the examination, the chapter collapses to a sequence that survives when detail fails. Identify the class from the mechanism, run the card for the agent in question, name the axis that separates the members, and state the class-defining hazard that constrains all of them. What consolidation buys is not more facts but faster and more reliable access to the facts already learned, so that under pressure the pharmacology arrives as a method rather than as a set of half-remembered associations about individual drugs.

> Within a class the mechanism is shared, so the choice is made from what differs: what else the agent binds, how fast it clears, and how much room lies between the exposure that helps and the exposure that harms.

Sources and further reading

Every dose, threshold, half-life and frequency in this chapter is grounded in one of the works below; where two sources set different figures, both are named in the text as a conflict rather than reconciled. Rating scales and monitoring instruments are described and cross-referred, never reproduced. Editions and publication years here were read from each work's own title and copyright pages rather than recalled.

Taylor DM, Barnes TRE, Young AH.
The Maudsley Prescribing Guidelines in Psychiatry. John Wiley and Sons. The 15th edition, first published 2025, is the chapter's principal prescribing source; the 14th edition, first published 2021, is retained where it carries a statement the 15th drops. The Maudsley Deprescribing Guidelines, first published 2024, is the source for withdrawal and tapering.

Sadock BJ, Sadock VA, Ruiz P.
Kaplan and Sadock's Comprehensive Textbook of Psychiatry. Eleventh edition, Wolters Kluwer, 2025. The principal reference across the chapter: receptor pharmacology, the antipsychotic and antidepressant classes, and the mood stabilisers. Its India chapter is cited separately where local prescribing practice is at issue.

Stahl SM.
Stahl's Essential Psychopharmacology: Neuroscientific Basis and Practical Applications, and the companion Prescriber's Guide, seventh edition, Cambridge University Press, 2021. Receptor mechanism and the profile-to-effect reasoning that organises this chapter.

Meyer JM, Stahl SM.
Stahl's The Clozapine Handbook, Cambridge University Press, first published 2020; and The Lithium Handbook, Cambridge University Press and Assessment, first published 2023. The two single-drug monographs behind the clozapine and lithium sections.

Brunton LL, Knollmann BC (Goodman and Gilman).
Goodman and Gilman's The Pharmacological Basis of Therapeutics, fourteenth edition. Receptor theory, dose-response and the therapeutic index as pharmacology rather than as prescribing. The publisher and year are deliberately not printed: the parsed copy held in this project has no title page and no copyright page, so neither could be read from the record, and neither is stated here from memory.

Tripathi KD.
Essentials of Medical Pharmacology. Eighth edition, Jaypee Brothers Medical Publishers, 2019. The pharmacology reference in widest use in Indian training, cited for class mechanism and for agents whose Indian use differs from the Western default.

Johnstone EC, Owens DGC, Lawrie SM, et al.
Companion to Psychiatric Studies. Eighth edition, Elsevier, Churchill Livingstone imprint, 2010. Class structure and the comparative pharmacology of the antidepressants.

Tasman A, Kay J, Lieberman JA, et al.
Tasman's Psychiatry. Fourth edition, John Wiley and Sons, first published 2015. Comparative within-class parameters, including the stimulant duration-of-action table.

Martin A, Bloch MH, Volkmar FR (Lewis).
Lewis's Child and Adolescent Psychiatry: A Comprehensive Textbook. Fifth edition, Wolters Kluwer, 2018. Cited for pharmacokinetic principle and for the stimulant pharmacology this chapter teaches at class level.

Indian Psychiatric Society.
Clinical Practice Guidelines. The source for the within-class Indian rankings of weight-gain and prolactin liability used in this chapter.

Also cited, in single passages.
Practical Psychopharmacology: Clinical Topics in Psychology and Psychiatry; the Shorter Oxford Textbook of Psychiatry; Stahl's Substance Use; and The Maudsley Prescribing Guidelines for Mental Health Conditions in Physical Illness. Each grounds one or two claims only, and each is named here by title alone because its edition and publisher were not read from the copy held in this project.

Index

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.