WEAVE SPRINT
PAPER IV · VOLUME SIX

MD and DNB psychiatry, theory revision

Neurology, Medicine and Recent Advances

Recent Advances and Landmark Studies

What changed recently, and which studies changed it. Read as evidence rather than as news, with the study behind each claim.

Dr. Wilfred D'souza · Dr. Niharika Reddy

WEAVE
CENTRE FOR INTEGRATIVE PSYCHIATRY

Weave
Centre for Integrative Psychiatry

WEAVE SPRINT
PSYCHIATRY OUTREACH PROGRAMME

Psychiatry that listens before it prescribes.

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





Assessments

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



Consultations and sessions

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



Family Library

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

`weave.clinic/family`

REFER, OR SEND SOMEONE OUR WAY

Your patient contacts Weave directly and decides whether to book.

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

CONTACT

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

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

WEAVE.CLINIC

Routine referrals only. This is not an emergency route.

WEAVE SPRINT · PAPER IV

Neurology, Medicine and Recent Advances

THE WEAVE SPRINT SERIES

TITLE

Weave Sprint, Paper IV, Volume Six: Recent Advances and Landmark Studies

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

WRITTEN BY

Dr. Wilfred D'souza and Dr. Niharika Reddy

PUBLISHED BY

Weave, Centre for Integrative Psychiatry · weave.clinic

EDITION

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

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

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

ACKNOWLEDGEMENT

A word of thanks

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

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

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

Dr. Wilfred D'souza

Foreword

This volume is one chapter, and the reason it stands alone is that it goes out of date faster than anything else in the set. Recent advances are usually revised from summaries and headlines. They are examined as evidence, which means the study, the design, the number of patients and what the finding does not show.

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

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

HOW TO USE THIS BOOK

The colour memory code

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

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

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

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

Contents

FRONT MATTER

Foreword 7

How to use this book 8

PART IX **Recent Advances and Landmark Studies** 11

Recent advances and landmark studies 11

Individual newer agents 111

Consolidate and apply 147

BACK MATTER

Subject index 160

About the authors 162

PART IX

IX

Recent Advances and Landmark Studies

What changed recently and which studies changed it, read as evidence rather than as news.

01 Recent advances and landmark studies

02 Individual newer agents

03 Consolidate and apply

PART 9 • Recent Advances and Landmark Studies

VISUAL 9.1



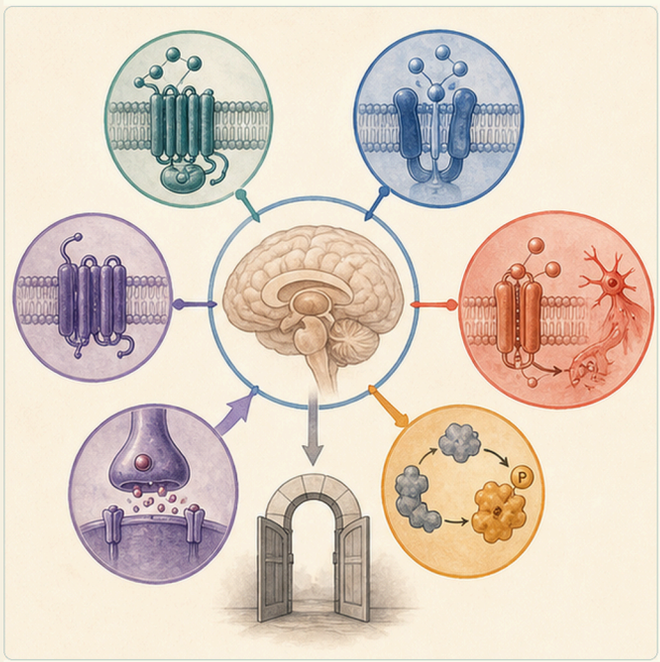
Rapid antidepressants use different treatment containers

- Glutamatergic
- Neurosteroid
- Monoamine
- Setting and monitoring
- Continuation plan

Rapid-acting antidepressant classes differ in mechanism, indication, delivery setting and continuation plan; speed alone does not make them interchangeable.

Source lock: day-38.final.html#sec-t04

VISUAL 9.2



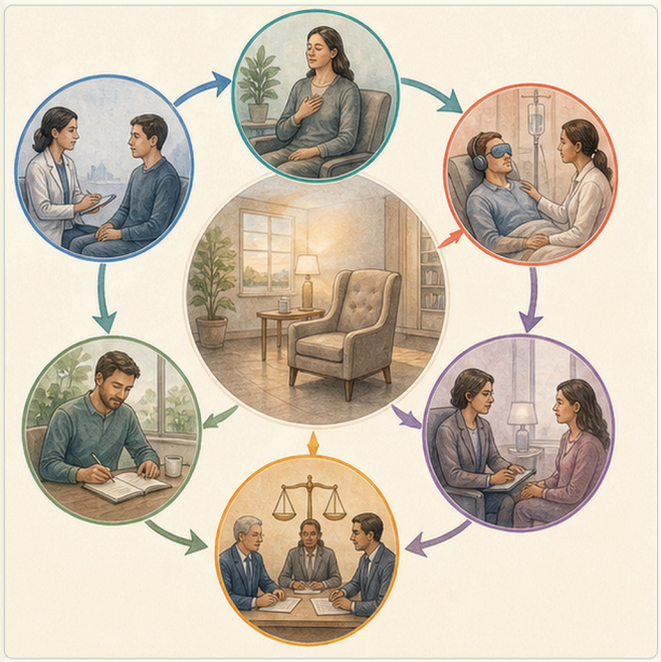
New mechanisms widen psychiatry beyond monoamines

- Muscarinic
- Orexin
- Sigma and NMDA
- Dopamine partial agonism
- PKC targeting
- Approval and indication

Recent agents expand mechanism classes, but each claim remains bounded by the medicine, indication, regulator and evidence base.

Source lock: day-38.final.html#sec-t13

VISUAL 9.3



Psychedelic treatment is more than a molecule

- Screening
- Preparation
- Supervised dosing
- Psychological support
- Integration
- Regulatory status

Psychedelic-assisted therapy is a structured intervention whose screening, support, setting, follow-up and regulatory status are part of the treatment.

Source lock: day-38.final.html#sec-t10

VISUAL 9.4



Computational psychiatry needs an evidence ladder

- Data
- Model
- External validation
- Bias and privacy
- Clinical integration
- Regulatory evidence

AI and computational tools require valid data, external testing, bias and privacy safeguards, and clinical integration before they can improve care.

Source lock: day-38.final.html#sec-t12

PAPER IV

Recent advances and landmark studies

Two bands . one complete notes document

Recent advances and landmark studies

1 · Recent advances and landmark studies for the psychiatrist: what this chapter owns, and what it points to

A drug is never simply approved: it is approved somewhere, for something, on a named authority, as of a date. This chapter gathers the agents, arguments and technologies that arrived in psychiatry most recently. The obvious risk in such a chapter is overlap, since most of these molecules are taught at depth elsewhere in these papers, and the boundaries later on this page handle it. The risk that actually costs marks is different, and installing it is this page's purpose: many of the questions this chapter can ask are about a status rather than a mechanism, and a status is not a property of a molecule.

That is a claim about most of the chapter and not about all of it, and the limit tells you where the discipline applies. Thirteen of the eighteen sections ahead have a regulatory record behind them, at the grain of one agent in one jurisdiction. Five do not, and not through oversight: an argument about a body of evidence, three technologies that are not medicines, and one landmark whose status is left deliberately to the chapter that owns it. In those five the examinable content is mechanism, appraisal and readiness, and a candidate looking for an approval to quote will find none and must not supply one. So the first move on any stem is to decide which of the two kinds of question you have.

1.1 The lens: a status claim is a four-field claim

A statement that an agent may be given, to whom and for what, is a claim about a regulator's document rather than about pharmacology, and this chapter treats it as a **four-field status claim**. The **four fields** are the jurisdiction, the exact indication or trial purpose in the source's own words, the authoritative source that carries it, and the date on which the status was true, which these notes record as **status true as of**. A sentence missing any one of the four is unmarkable, for a clinical reason rather than a stylistic one.

Jurisdiction	Indication	Source	Date
WHICH REGULATOR	IN THE SOURCE'S OWN WORDS	THE DOCUMENT, NAMED	WHEN IT WAS TRUE

Each field fails in a characteristic way. Drop the jurisdiction and the sentence is true somewhere and false where the patient is. Paraphrase the indication and the population quietly widens, because a licence is written for a defined group and the defining words are the ones that were reviewed. Leave the source unnamed and the claim rests on the speaker, which under examination means on memory. Leave the date off and a regulatory record is offered as a

permanent fact, when the reason this chapter exists is that these records moved recently and will move again.

RULE **An approval verb never travels alone.** Before writing that an agent is approved, indicated, authorised or licensed, put all four fields beside it: the jurisdiction, the indication in the source's own words, the document, and the date it was read. If any of the four is unavailable the sentence is not written shorter; it is written as a declared gap.

1.2 Jurisdiction is the hazard, and it does not run as a ladder

The organising fact of this chapter is that its agents are carried in up to **three jurisdictions** and that the cell, not the agent, is the unit of truth. Nothing about the three is ordered. The ledger behind these notes carries a grounded Indian row for lemborexant while its European cell is a declared gap on the centralised register, and a grounded Indian row for endoxifen while its United States cell is a declared gap after multi-route probing of the federal application database. Reading the United States first, the European Union next and India last as a ladder of maturity produces two confident wrong answers in this chapter alone.

The intuition behind the ladder is not stupid, which is why it survives. Volume does run that way: across the sixteen agents the ledger holds, the United States column carries the most grounded cells, the European Union fewer, India fewest. But a column total says nothing about any particular cell, and **regulatory recency** is the hazard this chapter names because a recent record is distributed unevenly, for reasons with little to do with maturity. An agent may reach one regulator first because that is where its sponsor filed, or reach India by a route that never went to a large market. Reasoning from column to cell is arithmetic where the question asked for a document.

TRAP **Filling an empty cell from a full one.** Asked about Indian availability, candidates answer from the American label, because the American record is the one they have read. The two cells are independently sourced, and this chapter holds pairs that point in opposite directions, so an inference from one to the other is not a cautious estimate. It is a fabricated citation wearing a real document's name. Say what the Indian record shows, or say what was searched and on whose authority a claim would rest.

1.3 Three states a cell can be in

Every cell reads in one of three ways, and knowing which one you have settles what may be written from it. A **declared gap** is the state most often misread, because it is a statement about an instrument and never about the world. Where the search ran on several routes, each positive-controlled in the same run, these notes call the result a **measured absence** rather than a probe silence, and even then one gap in this chapter turned out to be a probe miss once the record was retrieved another way. The third state is the one a candidate never expects: a **supportive source** that recites an approval without being one.

THE CELL READS	WHAT IT LICENSES	WHAT IT NEVER LICENSES
Grounded	The indication in the source's own words, for that jurisdiction alone, with its document and its date beside it	A sentence about any other jurisdiction, or the same sentence later without a fresh reading
Declared gap	What was searched, on which instrument, how sensitively, and whose authority a claim would ultimately rest on	Any categorical negative. Silence on an instrument is a fact about the instrument
Supportive source only	That a named body recited an approval, and the document in which it recited it	That the regulator granted it, or the recited date reused as an approval date

The middle state carries most of the weight, because this chapter has already had one gap turn out to be wrong: a cell that read as an absence was re-probed by another route and the record came back. So a declared gap is written as a description of a search rather than a verdict, on the terms the table above sets. A gap written that way can be closed by the next reader of the primary document. A gap written as a negative has to be retracted.

1.4 The instrument, and the five sections it cannot see

The shape of the apparatus behind these notes decides which questions it can answer. The record holds one row per agent per jurisdiction, sixteen agents across three regulators, every row carrying the same four fields. Those rows attach to thirteen of the chapter's eighteen sections. The five without them are the sections whose subject is not a licensed medicine: the serotonin hypothesis debate, an argument about a body of evidence rather than about any agent's licence; the artificial intelligence, biomarker and digital therapeutics sections, which are not what a drug-approval register records; and the anti-amyloid antibodies, where no row is written deliberately, so that none can be quoted from here.

The instruments behind this chapter are drug-approval instruments. When one returns nothing about a technology that is not a medicine, the silence is a fact about the instrument and not about the product, and a sentence converting it into a status is an invented claim wearing a real register's name. In those five sections the honest move is to say what the instruments cover and what they cannot see, then teach what is genuinely examinable: what the object is, what evidence would establish it, and what a service would have to know before using it.

IN INDIA

The Indian cell is the one most often filled in from memory, and the published CDSCO approval lists behind these notes carry a measured limit worth knowing before any silence in them is trusted: their positive controls all sit in documents dated well before the era this chapter covers, and one component substance of an agent taught here appears nowhere in that corpus under any spelling. So an Indian cell is either written from a document that was read, or declared as a gap naming what was searched and how sensitively. A brand on a pharmacy shelf, a CTRI trial registration, a no-objection certificate for a clinical trial and a committee minute recommending approval are none of them a marketing authorisation, and one agent here sits in exactly that state: a national committee's minutes recite an approval, and the row recording them licenses only the fact that they recite it.

1.5 A map of the chapter

With the lens in place the architecture becomes legible. The chapter runs in two declared halves: the recent advances of the last two years, then six individual newer agents held as whole-agent profiles. The map groups the eighteen sections into seven bands. It is not a list to memorise; it tells you which question a later section answers and which neighbour it borders.

BAND OF THE CHAPTER	THE SECTIONS IN IT	WHAT THE BAND IS FOR
The rapid-acting antidepressants	Esketamine and intranasal therapy; zuranolone and the neurosteroid antidepressants; rapid-acting antidepressants as a class	A comparative class map, so the category is not learned as a third ketamine monograph
Sleep-wake pharmacology	Dual orexin receptor antagonists and sodium oxybate	The sharpest regulatory contrast here: one class, four agents, three patterns across three jurisdictions
Mechanisms outside the established algorithms	Xanomeline-trospium and the muscarinic antipsychotics; endoxifen and the protein-kinase-C-targeted agents	Two mechanisms beside the established algorithms rather than inside them, and the most Indian section here
A landmark named and handed on	Anti-amyloid monoclonal antibodies	Why the class matters to psychiatry and not only to neurology, with its status delegated explicitly
Two arguments about evidence	The serotonin hypothesis debate; psychedelic-assisted therapy	What a body of evidence does and does not license, and how to appraise an intervention that is a package, not a molecule
Research to clinic	Artificial intelligence and machine learning; biomarker-driven and computational psychiatry; digital therapeutics and prescription apps	One readiness ladder read three times: analytical validity, clinical validity, clinical utility, regulatory authorisation, routine deployment
The individual newer agents	Asenapine; lurasidone; brexpiprazole; cariprazine; lumateperone; zuranolone as an agent profile	Whole-agent profiles read cell by cell, with the class primer given once instead of six times

Read the map both ways. Given a question, it says which band holds the answer and what kind of answer is wanted. Given a section, it says what its neighbours already cover, which stops an answer drifting into material being marked elsewhere.

1.6 What this chapter owns, and what it hands over

Naming the owning chapter is faster and better rewarded than rebuilding its material from memory. For the drug handovers below, that chapter holds the clinical pharmacology and this chapter holds only what changed, where, and on whose record. For the last two the neighbour holds the neurobiology and the service question instead, and this chapter holds no regulatory record for either, as the second paragraph above says. The handovers are not the whole picture, and treating them as though they were is its own error: this chapter owns seven of its eighteen sections in full, being the neurosteroid material and its agent profile, the muscarinic, endoxifen and lumateperone sections, and the artificial intelligence and biomarker sections. The research-to-clinic readiness ladder those last two build with the digital therapeutics section is this

chapter's as well. Where a section is owned in full there is no neighbour to point at, and an answer that gestures at another chapter has left the marks on the table.

- Ketamine and esketamine in treatment-resistant depression are taught in full in the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes, and the antidepressants and antipsychotics themselves in the Mood disorders: pharmacotherapy notes.
- The hypnotics and the oxybates belong to the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes and the Eating and sleep-wake disorders notes; the psychedelics as substances belong to the Substance use disorders notes.
- Anti-amyloid monoclonal antibodies are named here and taught in the Dementias and neurocognitive disorders notes, which own their status entirely; this chapter states none.
- Service-level deployment of anything digital belongs to the Community psychiatry and public health notes, and the serotonin hypothesis as neurobiology to the Mood disorders: classification, depressive and bipolar syndromes, neurobiology notes and the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes.

PITFALL

Quoting a status you learned once as though it were still true. Every row behind this chapter carries the date its status was true, because this is the most recently changed material in the syllabus: an indication is widened, a second regulator authorises the same agent for something narrower, a product is discontinued while its approval stands. A candidate who memorised a status and reproduces it much later has not made a small error of currency: the sentence carries no date, so it offers a stale reading as a permanent fact. Learn the four fields and the shape of each record rather than the record itself, and say when you last read it. In the clinic it means checking the document before prescribing a recently licensed agent, not after a query is raised.

1.7 What the examination rewards

A small number of moves earn most of the marks here, and they recur across every band in different clothing. Naming them before the detail arrives is worth more than any single fact, because the detail in this chapter is unusually perishable and the moves are not.

- Answering at the cell rather than at the agent, so a question about one jurisdiction is never answered out of another's record.
- Producing the four fields as a sentence, instead of a verdict with the qualification left to be assumed.
- Reading a silence as a statement about an instrument, and saying what was searched, how sensitively, and on whose authority a claim would rest.
- Giving an ordering with its basis named wherever no quantity is available, rather than a figure that would have to be invented to be given.
- Deciding whether the stem wants a status or a construct, then handing over what a neighbour owns in one sentence rather than rebuilding it.

Those five are also, in the same order, the five ways an answer here goes wrong, which is why they are habits rather than advice. A candidate who runs them will sometimes have to say that a record was searched and not found, or that an ordering is all the evidence supports. Both are complete answers here, because the alternative is a confident sentence nobody can source.

GOOD TO REMEMBER

Before recommending, prescribing or teaching any agent from this chapter, ask four questions in order. Which regulator's document supports this use in the country where this patient is being treated. What does that document say the indication is, in its own words rather than in yours. When did you last read it. And is the thing you are about to say a status claim, needing all four fields, or a mechanism claim, needing none of them and needing a stated limit instead. A clinician who answers those four before speaking will not be the person who tells a family a treatment is available on the strength of a label written for another country.

1.8 Holding it together

Under examination pressure the detail goes first, because a date and a verbatim indication are harder to hold than a mechanism. What has to survive is a routing sequence. The scaffold below is a teaching device for that sequence rather than a validated instrument, and it will structure an answer on any of the eighteen sections.

MNEMONIC

Kind, Currency, Grain, Owner

Kind: which of the chapter's three kinds of section the stem sits in, an agent with a regulatory record, an argument about a body of evidence, or a technology on a readiness ladder. **Currency**: what the answer is paid in, a status with its four fields, or a construct with its mechanism, its evidence and its limit. **Grain**: if it is a status, answer at one agent in one jurisdiction, and say which of the three states that cell is in. **Owner**: who holds the clinical pharmacology, which for most sections is a neighbouring chapter named in a sentence, and for seven of the eighteen is this chapter, where a full account is owed and a cross-reference earns nothing.

For a short answer, run the four and stop. For a long essay, expand each into a paragraph and take the specifics from the section that owns them. In a viva, make the sequence audible: name the jurisdiction, say what the document says and when it was true, and say what you would have to read to answer for anywhere else. That last sentence separates a candidate who has understood this chapter from one who has memorised it, because it concedes a limit without conceding the question.

Every status sentence carries its jurisdiction, its exact indication, its authoritative source and the date it was true, and no status is ever carried across a border.

2 · Esketamine and intranasal therapy, including monotherapy indication updates

An approval is not a property of a molecule; it is a sentence a named regulator wrote in a named country on a named date. Intranasal esketamine is where that discipline is examined, and the marks are lost in one move: a status carried across a border, almost always from the United States inward, into a sentence true somewhere and false where the patient is. Delivery, dissociation and blood-pressure monitoring, the observation period and the access problem belong to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes and are not repeated here.



Fig 38.1 **The four-field status ledger applied to esketamine, one record per regulator.** The United States record, read 3 August 2026, permits the spray alone for its resistance indication, as monotherapy or in conjunction with an oral antidepressant. The European record, read 1 August 2026, carries two authorised indications, one requiring a SSRI or SNRI and the other oral antidepressant therapy, so neither authorises it alone. The Indian record is a declared gap: no CDSCO marketing approval record located, across 66 published approval lists, with the document that would close it named. The monotherapy question therefore has three answers and no single one. (openFDA Drugs@FDA NDA211243; the European Annex I section 4.1 for Spravato; the CDSCO approval lists and the Indian trials registry.)

2.1 Four fields, and the sentence that fails without them

A regulatory status is usable only when it carries four fields together: the jurisdiction, the exact indication in the words the source itself uses, the authoritative source, and the date on which the status was true. "Esketamine is approved for depression" fails on all four, and last on the date, in a domain where label text is amended and last year's correct answer can already be wrong.

Every statement below was read from its primary source on 1 August 2026, with the European indication text re-read verbatim on 2 August 2026 and the United States label text re-retrieved on 3 August 2026, and each is offered as true as of that reading and nothing more.

A status never travels, and the source record says so: the label wording is not a statement about any other jurisdiction and must never be carried across one. And "approved internationally" is not an answer: name the regulator, or say that no record was located and what was searched.

RULE **A status sentence is written in one jurisdiction and stays there.** An American indication does not imply a European one, and neither implies anything about a country whose regulator was never asked. Where the record for a jurisdiction is empty, say the record is empty, in that jurisdiction, on that date.

Marketing authorisation is a permission granted by a national regulator, and it is not availability. Expert committee minutes, trial no-objection certificates, registry entries, brand listings and pharmacy availability are none of them marketing authorisation.

2.2 The United States: the indication that carries monotherapy

Two indications sit on the current United States label. The first is for the treatment of treatment-resistant depression (TRD) in adults as monotherapy or in conjunction with an oral antidepressant. The second is for depressive symptoms in adults with major depressive disorder (MDD) with acute suicidal ideation or behavior in conjunction with an oral antidepressant. The resistance indication offers a choice, **monotherapy** or combination; the suicidality indication offers none, the oral antidepressant being part of it rather than an option inside it.

The label then attaches three Limitations of Use. The first: the effectiveness of SPRAVATO in preventing suicide or in reducing suicidal ideation or behavior has not been demonstrated. So the licence is for depressive symptoms in a patient who has **acute suicidal ideation or behavior**, and it is not a licence for treating the suicidality; sliding from the first formulation to the second claims an effect the label explicitly declines. The second: use of SPRAVATO does not preclude the need for hospitalization if clinically warranted, even if patients experience improvement after an initial dose of SPRAVATO. A rapid response is not a disposition decision. The third: SPRAVATO is not approved as an anesthetic agent. The safety and effectiveness of SPRAVATO as an anesthetic agent have not been established.

The status was read from the openFDA Drugs@FDA interface, application NDA211243, with the current structured product label effective 13 March 2026, on 3 August 2026. The record carries a caution about itself: openFDA is a mirror of Drugs@FDA and of the label archive operated by the agency, and not the Drugs@FDA application itself.

The earliest approval entry in that application record is dated **5 March 2019**, and the field is named for what it holds, the earliest approval-status date in the record, not for what a reader would like it to mean. The record carries no date for the monotherapy wording: the label in force at the reading contains it, and nothing in the source read here says when it entered. So the defensible sentence is that the United States indication included monotherapy on the label effective 13 March 2026, read on 3 August 2026; the indefensible one names the year monotherapy was added.

GOOD TO REMEMBER

An indication that permits monotherapy permits it; it does not recommend it, and it does not make the evidence for the drug alone equal to the evidence for the drug alongside an antidepressant. The maintenance trial and the head-to-head comparison taught in the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes are both add-on designs there, so a durability or comparator figure from either carries the design it came from and does not attach to a monotherapy sentence it was never generated under.

2.3 The European Union: two authorised indications, and neither is a monotherapy

The European authorisation covers the same product and is not the same licence. Annex I of the **Summary of Product Characteristics**, section 4.1, carries two therapeutic indications, and both must be reproduced before the licence has been described. The first: Spravato, in combination with a SSRI or SNRI, is indicated for adults with treatment-resistant Major Depressive Disorder, who have not responded to at least two different treatments with antidepressants in the current moderate to severe depressive episode. The second: Spravato, co-administered with oral antidepressant therapy, is indicated in adults with a moderate to severe episode of Major Depressive Disorder, as acute short-term treatment, for the rapid reduction of depressive symptoms, which according to clinical judgement constitute a psychiatric emergency.

Neither indication authorises the spray by itself: the first requires an SSRI or SNRI, the second requires oral antidepressant therapy, so there is no European monotherapy indication of any kind. And the sentence that European esketamine is indicated only for treatment-resistant depression is false, because the second indication is not about resistance. It is an acute, short-term indication for what clinical judgement calls a **psychiatric emergency** inside a moderate to severe episode, and a patient can meet it at first presentation, having failed nothing.

That the second indication is really there was tested. The European row was checked against the primary document on three instruments: whether its text is verbatim in the product information, whether it is the primary text cut short and closed with a full stop so that nothing signals the loss, and whether the count of indications carried matches the count in the document. Two indications are in the document and two are on the row. A one-indication summary of a two-indication licence reads as complete, which is why the third instrument counts.

Spravato holds product number EMEA/H/C/004535, its marketing authorisation holder is Janssen-Cilag International NV, and its status on the European Medicines Agency register is Authorised. The initial European Union marketing authorisation, by European Commission

decision, is dated **18 December 2019**. The most recent European Commission decision recorded against the product is dated **12 December 2024**, and that date needs a warning: the field records the latest decision and does not state what it did. Nothing read here says it touched the indication, and converting a decision date into an indication change manufactures the change out of a timestamp.

One limit belongs on every European sentence here. This is an authorisation for the named product through the **centralised procedure**, and it says nothing about other products containing the same active substance that may be authorised nationally by the decentralised or mutual-recognition procedures, which the central register does not cover by construction. "Authorised in the European Union" therefore means authorised centrally, for this product, and an answer should say so.

2.4 The monotherapy line, drawn regulator by regulator

The line falls differently in each jurisdiction.

JURISDICTION	THE RESISTANCE INDICATION, IN THE SOURCE'S OWN TERMS	MAY THE SPRAY BE GIVEN ALONE?	NAMED PARTNER DRUG
United States	Treatment-resistant depression in adults, as monotherapy or in conjunction with an oral antidepressant	Yes, expressly, for this indication	An oral antidepressant, no class named
European Union	Treatment-resistant major depressive disorder in adults who have not responded to at least two different treatments with antidepressants in the current moderate to severe depressive episode	No	A SSRI or SNRI, named and restricted
India	No CDSCO marketing approval record located as of 1 August 2026	The question has no answer in the record	The question has no answer in the record

One difference is not in the table. Only the European indication writes its own entry criterion, counting failed antidepressant treatments within the current episode; the American resistance indication names the condition and leaves the counting elsewhere.

Both second indications require an oral antidepressant alongside, and there the agreement stops: the American is anchored to a symptom on which the label then declines to claim an effect, the European to a clinical judgement about the situation as a whole.

1 of 2 US INDICATIONS PERMITTING MONOTHERAPY	**2** AUTHORISED EU INDICATIONS, BOTH COMBINATION	**66** CDSCO APPROVAL LISTS SEARCHED	**0** INDIAN APPROVAL RECORDS LOCATED

TRAP **Writing "esketamine is not a monotherapy" as a fact about the drug.** It is a fact about a jurisdiction. The add-on-only sentence matches the European wording, where both authorised indications require a partner antidepressant, and it does not match the current United States label, where the resistance indication reads as monotherapy or in conjunction with an oral antidepressant. It is also the sentence carried in the Mood disorders: pharmacotherapy notes, and that adjunctive position holds on the European wording; the American half of it is what this dated update revises. The mirror-image error is just as costly: writing that esketamine is approved as a monotherapy, full stop, states an American permission as though it were a global one.

TRAP **Writing that the European licence requires an SSRI or SNRI.** True of the resistance indication and untrue of the other one. The SSRI or SNRI restriction is printed in the treatment-resistant indication; the psychiatric-emergency indication says co-administered with oral antidepressant therapy, the wider phrase. A candidate who has learned only the resistance indication will state the narrow partner rule for a patient it was never written for. Quote the indication you mean, and quote it whole.

2.5 India: what was searched, what was not found, and what would close it

The Indian position is the one an Indian examination cares about most and the one most often invented. Stated exactly: no CDSCO marketing approval record was located, and esketamine is absent from all 66 published CDSCO approval lists, as of 1 August 2026. That is a statement about a search, which has to be described before it means anything.

Three instruments were run. The first was the corpus of 66 published CDSCO approval lists, every one of which yielded an extractable text layer, so no absence in it is an artefact of an unreadable scan. The second was a search of the regulator's own domain with each returned document downloaded and its text verified rather than its title trusted. The third was a search of the national clinical trials registry domain, validated against a positive control before any zero result was believed. The design point: a search never shown to find what it should find cannot prove that something is missing.

Five documents came back: four sets of subject expert committee minutes for neurology and psychiatry, and one international cell file. All four sets of minutes were fetched, all carry an extractable text layer, and none of them names esketamine. Those **subject expert committee minutes** were search-index false positives, not evidence. No indexed trial record was located on the national registry either.

The corpus is not empty of the parent molecule, which is where a careless reader manufactures an approval. It contains ketamine hydrochloride as an anaesthetic agent approved in August 1971, and topical ketamine combinations from 2010. Neither of those is esketamine. Neither a racemic anaesthetic licensed for another purpose nor a topical combination licensed for another route is a marketing authorisation for an intranasal antidepressant; and the American label states in terms that SPRAVATO is not approved as an anesthetic agent.

What would close the gap is nameable, which is the test of an honest one. The missing document is an affirmative CDSCO or DCGI new-drug permission, or a current CDSCO-approved label, or

an entry in an official CDSCO approval list. Nothing else substitutes. Expert committee minutes are supportive at best and are not themselves marketing authorisation; a clinical-trial no-objection certificate, a written confirmation from the international cell, a CTRI registration, a brand listing, pharmacy availability and manufacturer copy are none of them marketing authorisation either. Name the document that would settle it, and say whether you have it.

IN INDIA

The sentence to write, in an examination and in a clinical note, is that no CDSCO approval record for intranasal esketamine was located as of August 2026, followed by what was searched: the published approval lists, the regulator's own domain, and the trial registry. Do not write that esketamine is unapproved in India, which is stronger than the evidence supports, and do not write that it is approved, because nothing retrieved says so. A brand name, a price or a distributor's letter is not that document. For a patient or a family: this treatment has a licence elsewhere and no located Indian licence, and what a service can actually offer today is the question to answer first.

2.6 Why the Indian finding is a declared gap and not a negative

A **declared gap** states that a search of named instruments did not retrieve a document, and names the document it wanted. A negative states that the thing does not exist. The first is defensible on the evidence in hand; the second needs an instrument whose silence means something, and that sensitivity has to be measured, not assumed.

Here it was measured, and it is poor in exactly this chapter's window. Every drug the corpus is known to find sits in a document dated 2023 or earlier, while every agent in this chapter is a recent approval, and independent sensitivity in the 2024 to 2026 window is 0 of 1. The corpus is therefore no warrant for a gap on a recently approved agent: it has never been shown to find one.

The stronger leg is that the regulator's own documents contradict its own lists. CDSCO subject expert committee minutes of 16 April 2025 state that a lumateperone capsule had already been approved in the country on 26 December 2024, yet lumateperone appears in none of the 66 lists, and the list that declares full-year 2024 coverage in its own running header carries entries dated 17 and 30 December 2024, so it is not truncated before the approval those minutes recite. A list that misses an approval its own committee recites, inside its own declared coverage window, is a list whose silence cannot be read as absence.

Write the conclusion narrowly. Absence from this corpus is evidence and not proof, and the claim being made is that no CDSCO approval record was located, not that the drug is unapproved in India. Note what was deliberately not done: the absence was established from Indian sources only, and no American or European value was consulted to produce any part of it. A jurisdiction's emptiness is established in that jurisdiction, never inferred from another's fullness.

PITFALL

Counting search hits instead of reading them. The regulator-domain search returned five documents and none names esketamine; a count would have recorded five findings where the correct number is zero. The error runs the other way in the trial registry, where a zero was believable only because a **positive control** had returned a known record first. Open what the search returns, and prove the instrument can see before reporting that it saw nothing.

2.7 Reading a regulatory record without being fooled by it

Three failure modes turned up while this position was assembled, each producing a confident, well-cited, wrong sentence.

- **Corroboration that is circular.** The European indication was first taken from the agency's published medicines register and checked against the public assessment report page. The register is a machine-generated dump of those same pages, so the corroboration was circular and blind to precisely the failure that had occurred, a stale summary layer. Annex I section 4.1 of the Summary of Product Characteristics is the genuinely independent instrument and is the authority for the indication text; the register and the report page corroborate product identity and authorisation status only.
- **Truncation that closes with a full stop.** A summary reproducing the first part of an indication and stopping at a sentence boundary reads as complete, because nothing signals the loss. The European row was tested for it from the start; the American row was not, and the stored copy of this very label stopped mid-clause at "does not preclude the need ", hiding a whole third limitation until an author refused to complete it from memory. The instrument was then pointed at the American column too, and the repaired row passed at 1649 characters against 1649. Quote an indication to its end, and where the copy in hand is cut off, say so and quote only the part that is whole.
- **A date wearing the wrong field name.** The register's decision-date column carries the latest European Commission decision, not the first; the build proved it on another product in the same workbook, where the column showed a 2026 date against an authorisation granted in 2014. Label a date for what it is: initial marketing authorisation, latest decision, earliest approval entry in an application record. A true date of one thing published under the field name of another is factually sourced and substantively false.

All three share one shape: the instrument returned something, the reader accepted it, and nobody asked what it was incapable of showing. The discipline is not what did the source say, but what could it not have told me.

2.8 What a complete answer contains

Reduced to four sentences:

- United States, on the label effective 13 March 2026, read 3 August 2026: treatment-resistant depression in adults, as monotherapy or in conjunction with an oral antidepressant; and depressive symptoms in major depressive disorder with acute suicidal ideation or behavior, in

conjunction with an oral antidepressant, under the three Limitations of Use quoted whole above.

- European Union, on the Summary of Product Characteristics read 1 August 2026 and re-read verbatim the next day: two authorised indications, treatment-resistant major depressive disorder in combination with an SSRI or SNRI, and the acute short-term reduction of depressive symptoms constituting a psychiatric emergency, co-administered with oral antidepressant therapy. Neither is a monotherapy indication.
- India, as of 1 August 2026: no CDSCO marketing approval record located, across 66 published approval lists, a verified regulator-domain search and a positive-controlled trial-registry search. A declared gap, not a statement that the drug is unapproved.
- The monotherapy question therefore has three answers and no single one.

MNEMONIC

Where, What, Who says so, When

The four fields of a status claim, in the order an examiner checks them. Where is the jurisdiction; What is the indication in the source's own words, whole and to its end; Who says so is the named regulator and the named document, never "approved internationally"; When is the date the status was read. An answer missing the fourth is an undated assertion in a domain that changes.

Esketamine permits monotherapy in one jurisdiction, forbids it in another, and has no located record in a third: the monotherapy question has no answer until a regulator and a date are named.

3 · Zuranolone and neurosteroid antidepressants for postpartum and major depression

This is the class that turned an endogenous hormone into a fourteen-day prescription, and then discovered how narrow the licence for it would be. Neuroactive steroids are the first antidepressants reasoned backwards from obstetric physiology rather than forwards from a monoamine, and two reached a regulator: brexanolone, an infusion now marketing status Discontinued in the United States, and zuranolone, the oral successor. The marks sit in three places: the subunit, the clock and the border.



Fig 38.2 Zuranolone and the neurosteroid line from brexanolone, marketing status Discontinued: the mechanism, the two evidence bases read separately, and the position that follows. What carried forward from the Discontinued infusion is a single course of treatment with no maintenance limb, and the United States schedule; what did not is the monitoring service it required. Each indications section names one indication: postpartum depression (PPD) in adults in the United States, and postpartum depression (PPD) in adults following childbirth (see section 5.1) in the European Union. (Zurzuvae EMA EPAR Annex I SmPC; ZURZUVAE and ZULRESSO United States labelling, the latter marketing status Discontinued; the EU medicines register; the published CDSCO approval lists.)

3.1 Why a steroid, and why after delivery

Allopregnanolone is a progesterone metabolite and an endogenous neuroactive steroid modulating inhibitory transmission. The United States label for the intravenous agent: brexanolone is a neuroactive steroid GABA-A receptor positive modulator chemically identical to endogenous allopregnanolone, marketing status Discontinued. Zuranolone is the synthetic relative, not the identical molecule. Its European Summary of Product Characteristics calls it a **neuroactive steroid** showing potent **positive allosteric modulation** of the GABA-A receptor, enhancing GABAergic activity at **synaptic and extrasynaptic GABA-A receptors** and increasing their cell surface expression in vitro.

That last clause is the examinable difference. The benzodiazepine site needs a gamma subunit, so benzodiazepines act at synaptic receptors and are essentially inactive at the delta-containing extrasynaptic receptors carrying the persistent background current. Neuroactive steroids bind elsewhere and modulate both. The United States brexanolone label, marketing status Discontinued, shows it: the drug potentiated GABA-mediated currents from recombinant human receptors expressing alpha1 beta2 gamma2 subunits, the classic synaptic assembly, and alpha4 beta3 delta and alpha6 beta3 delta subunits, which are extrasynaptic. Modulating the extrasynaptic population modulates **tonic inhibition**, the standing tone rather than the phasic response. An answer saying only "it works on GABA" has given away the subunit, the part being tested.

The physiology that pointed the class at the puerperium is a hypothesis. Allopregnanolone rises through pregnancy with placental progesterone and falls abruptly at delivery; the proposal is that in a susceptible subgroup the GABA-A system fails to re-adapt. Both regulators decline to endorse it. The Discontinued intravenous agent's label says the mechanism of brexanolone in postpartum depression is not fully understood but is thought to relate to positive allosteric modulation of GABA-A receptors, and the European zuranolone document says only that it may exert antidepressant effects by enhancing GABAergic inhibition. Print the source's own hedge.

RULE **The mechanism is why this is a course and not a regimen.** The United States licence for brexanolone, Drugs@FDA NDA211371, approved for postpartum depression on 17 June 2019 and now marketing status Discontinued, administers one 60 hour infusion. The European licence for zuranolone, Zurzuvae, EMEA/H/C/006488, authorised for postpartum depression on 17 September 2025, sets a single 14 day course. Neither document carries a maintenance schedule, which is why every question about continuation, relapse and a second course meets an absence of evidence, not a dosing table.

3.2 Brexanolone, the Discontinued first-in-class, and its delivery model

Brexanolone, marketing status Discontinued, proved the class could work. Its United States labelling indicates it for postpartum depression in patients 15 years and older. Administration is a **continuous intravenous infusion over 60 hours**: 30 micrograms per kilogram per hour for the first 4 hours, 60 for the next 20 hours, 90 for the following 28 hours, then 60 for 4 hours and 30 for the last 4 hours, holding at 60 through the middle block if 90 is not tolerated. The shape: open low, hold high, taper, never stopped abruptly.

The monitoring turned it into a service. This Discontinued product's label carries a boxed warning for excessive sedation and sudden loss of consciousness, and requires continuous pulse oximetry with an alarm, sedation assessment every 2 hours during planned non-sleep periods, and that patients be accompanied during interactions with their child or children while the infusion runs. Distribution ran through a **Risk Evaluation and Mitigation Strategy**, and the product is Schedule IV under the Controlled Substances Act. In the trials behind this Discontinued agent's licence, sedation or somnolence forced dose interruption or reduction in 5 per cent of treated patients against none on placebo, loss or altered state of consciousness occurred in 4 per cent against none on placebo, and recovery took 15 to 60 minutes. Most services cannot deliver that whatever the effect size, which is the case for an oral successor.

TRAP **A true date wearing the wrong field name.** The United States date of approval for brexanolone, marketing status Discontinued, is **17 June 2019**, while the regulator's submission-status field for the original application holds 19 March 2019. Both are real, and the approval letter explains the gap: because controlled-substance scheduling was incomplete, it states the date of approval shall be the date the Drug Enforcement Administration publishes a Federal Register notice announcing interim final scheduling. A date is not a fact until you know which event it timestamps.

3.3 Zuranolone: the shape of a fourteen-day treatment

The European posology sets **50 mg once daily for 14 days** as a **single course of treatment**, reducible to 40 mg once daily for 14 days if not tolerated, and states dosing may be stopped without down-titration. There is no maintenance limb and no second course: the same document records that no data on additional treatment in the case of a relapse or an insufficient response to zuranolone are available.

The capsules are taken in the evening with fat-containing food, as either a meal or a snack, an exposure requirement, since absorption is food-dependent. A missed dose is not made up. Zuranolone may be used alone or with stable background oral antidepressant therapy, unlike agents licensed only as add-ons.

Zuranolone is metabolised primarily by CYP3A. Strong inhibition raises exposure, so the dose drops to 30 mg once daily through the 14 day treatment period when a strong CYP3A inhibitor is co-prescribed, and grapefruit is avoided for the same reason. Induction does the opposite: exposure is reduced by 85 per cent in the presence of rifampicin, which is why inducers are avoided outright rather than dose-adjusted. In an Indian setting that list is not hypothetical: carbamazepine, phenytoin, phenobarbital and rifampicin. Renal and hepatic adjustment and the full adverse-reaction table sit in the agent-reference entry.

3.4 The postpartum evidence, read the way an examiner reads it

Two development programmes, one population, and no licence to mix their numbers.

PROGRAMME	WHO WAS STUDIED	PRIMARY ENDPOINT	WHAT IT SHOWED
Brexanolone, status Discontinued: two placebo-controlled studies, Studies 1 and 2	Women aged 18 to 45 meeting DSM-IV criteria for a major depressive episode, onset in the third trimester or within 4 weeks of delivery; Study 1 severe, HAM-D total at or above 26, Study 2 moderate, HAM-D total between 20 and 25	Change from baseline in HAM-D total at the end of infusion, the sixtieth hour; the thirtieth day pre-specified as secondary	Severe illness: the 90 microgram per kilogram per hour target beat placebo by 3.7 points and the 60 microgram target by 5.5; moderate illness: the 90 microgram target by 2.5
Zuranolone: one randomised, double-blind, placebo-controlled multicentre study, 217-PPD-301, NCT04442503, SKYLARK	Women aged 18 to 45 meeting DSM-5 criteria for a major depressive episode, baseline HAMD-17 at or above 26, onset limited to the third trimester or within 4 weeks of delivery, treatment started up to 12 months after childbirth, followed 4 weeks after the 14 day course	Change from baseline in HAMD-17 at the fifteenth day, the day after the course ends	Least-squares mean change minus 15.6 against minus 11.6, a treatment difference, 4.0 points, 95 per cent confidence interval 1.7 to 6.3, p equals 0.0007

Speed: separation is present by the third day, at 3.4 points, the property separating this class from a conventional antidepressant. Durability inside the trial: the difference is still measurable at the forty-fifth day, at 3.5 points, a month after the last capsule, where follow-up stops. The placebo arm improved by nearly 12 points, so the drug effect is a modest increment on a large non-specific recovery, as a naturally remitting condition with intensive trial contact produces.

4.0	57.0%	38.9%	14
TREATMENT DIFFERENCE IN HAMD-17 POINTS	RESPONDERS ON ZURANOLONE	RESPONDERS ON PLACEBO	DAYS, THE WHOLE COURSE

Response, at least a halving of the baseline total, was reached by 57.0 per cent against 38.9 per cent on placebo, a relative risk 1.409 and an absolute difference, 16.4 percentage points. Among those responders, three subjects, 5.7 per cent, relapsed by the trial's definition of two consecutive totals at or above 20 after the fifteenth day, and none met the definition of rebound. Two cautions: the denominator is small, so the confidence around 5.7 per cent is wide, and the window closes four weeks after the course.

3.5 What happened to the major depressive disorder claim

Most answers go wrong here by writing a sentence that is nearly true. Both regulatory records held for zuranolone name one indication, and it is not major depressive disorder. The United States labelling reads that ZURZUVAE is indicated for the treatment of postpartum depression (PPD) in adults, carried twice in that record. The European therapeutic-indications section, Annex I section 4.1, reads that Zurzuvae is indicated for the treatment of postpartum depression

(PPD) in adults following childbirth (see section 5.1). Both are quoted as printed, parentheses included. Neither indications section carries a second indication.

Set that against the trial population. Every patient in the pivotal study met criteria for a **major depressive episode**; what made the study postpartum was the clock, onset restricted to the third trimester or the 4 weeks after delivery. The untested step is from that onset window to no onset window at all.

Three clocks run through this topic; merging them produces a wrong sentence. The DSM-5 peripartum-onset specifier covers onset during pregnancy or in the 4 weeks after delivery. The trial's onset window is narrower antenatally, admitting only the third trimester. The treatment window is wider, since subjects could start treatment up to 12 months after childbirth provided symptoms began inside the narrow window. A woman presenting at nine months with an episode that began at eight matches none of the three, and a stem built around her is deliberate.

TRAP **Reading the postpartum evidence as depression evidence.** This chapter holds no trial evidence, in any jurisdiction, for a neurosteroid in major depressive disorder outside the peripartum window, and no regulatory record of an application, opinion or decision on such an indication from the United States Food and Drug Administration, the European Medicines Agency or the Central Drugs Standard Control Organisation. That is a declared gap in this chapter's evidence, not a finding that such a programme did or did not exist. Writing "zuranolone is a new rapid-acting oral antidepressant for depression" crosses the indication boundary and the evidence boundary at once.

One further argument comes from the shape of the licence, not from a decision this chapter lacks: the whole authorised European exposure is a single fourteen-day course, with no data on treatment after relapse. Major depression outside the puerperium is not defined by an event that ends, so a treatment with no second course and no maintenance limb has a structural problem in that indication whatever its acute effect.

3.6 Sedation, the driving rule, and the abuse question

Every frequency here is from the European Summary of Product Characteristics. The most frequent adverse reactions were somnolence at 27.6 per cent, dizziness at 13.3 per cent and sedation at 11.2 per cent; the reaction classed as serious was confusional state, at 1.3 per cent. Only 2 per cent discontinued for an adverse reaction, while 14.3 per cent had a dose reduction or interruption. Among those on the higher dose, 69.7 per cent of somnolence events and 58.3 per cent of sedation events occurred within the first 2 days of treatment. Warn her that the first two evenings are the heaviest.

The same European document counsels patients not to drive or engage in other potentially hazardous activities until **at least 12 hours after taking each dose**, and advises they may not be able to judge their own ability; the United States labelling carries the same restriction in a boxed warning. Two driving-simulation studies found dose-dependent impairment nine hours after dosing, and modelling of the lateral-position data put median placebo-corrected impairment

below the threshold for a 0.05 per cent blood alcohol concentration only by the twelve-hour mark.

PITFALL

Treating the sedation as a side effect to be tolerated rather than as the thing structuring the prescription. The patient is by definition caring for an infant, and the two documents handle that differently. The United States label of the Discontinued intravenous product requires a companion during interactions with the children. The European zuranolone document names driving and hazardous activities and does not name infant care among them, so an answer must not attribute to zuranolone a chaperone requirement the European document does not name. The counselling point: she may not judge her own alertness reliably, and night-time infant care over the first evenings is a conversation to have before the first capsule.

Abuse potential is assessed in that same European document. In a human abuse potential study in 60 recreational users of central nervous system depressants, zuranolone at 30, 60 and 90 mg showed dose-dependent abuse potential against alprazolam at 1.5 mg and 3 mg on positive subjective measures including drug liking and good drug effects, while clinical-trial data indicate low physical dependence potential. Liking is not dependence. Caution is advised where there is a history of abuse of or addiction to alcohol or other substances.

3.7 Pregnancy, contraception and breastfeeding

In the European document zuranolone is contraindicated during pregnancy, on animal reproductive toxicity and an absence of human data, and women of childbearing potential must use effective contraception during treatment and for 7 days following discontinuation. Contraception is part of the prescription, and cover runs a week past the last capsule.

The numbers and the instruction pull against each other; print both. A lactation study measured the **relative infant dose** under 1 per cent: the mean was 0.357 per cent at a 30 mg maternal dose, and simulation put the expected mean at 0.738 per cent and 0.984 per cent for a 50 mg maternal dose at two assumed milk intakes. Despite figures that low, the European document instructs that breastfeeding should be discontinued during treatment unless, in the judgement of the healthcare professional, the benefits of breastfeeding outweigh the possible risks to the child, because the effect on a breastfed infant is unknown and milk-production data are limited. The United States label for the Discontinued intravenous agent records a higher relative infant dose, 1 to 2 per cent of the maternal weight-adjusted dosage, with oral bioavailability under 5 per cent keeping infant exposure low. Quantified low exposure and a precautionary instruction are not a contradiction.

The age boundaries differ and must not be merged. The European document records that for zuranolone, safety and efficacy in postpubertal females under 18 years old have not been established, and there is no relevant use in prepubertal females. For the Discontinued intravenous agent the United States indication runs to patients 15 years and older.

GOOD TO REMEMBER

Two questions decide relevance, and neither is pharmacological: when the episode began, and whether she is breastfeeding. The first decides whether any evidence held here applies to her; the second, whether the licensed instruction can be followed without imposing a loss she may not accept.

3.8 Status by jurisdiction: what each regulator's record actually says

Every regulatory sentence here carries four fields, and one missing a field is not a status sentence: the jurisdiction, the exact indication or purpose, the authority the claim rests on, and the date the status was true. The interesting rows are the empty ones.

AGENT AND JURISDICTION	WHAT THE RECORD SAYS	AUTHORITY RELIED ON	STATUS TRUE AS OF
Zuranolone, United States	Indicated for the treatment of postpartum depression (PPD) in adults. Earliest approval-status date in the record, 4 August 2023	Drugs@FDA, NDA217369	1 August 2026
Zuranolone, European Union	Indicated for the treatment of postpartum depression (PPD) in adults following childbirth (see section 5.1); authorised as Zurzuvae, product number EMEA/H/C/006488, holder Biogen Netherlands B.V.; **marketing authorisation** valid throughout the European Union from **17 September 2025**	Indication from Annex I section 4.1 of the Summary of Product Characteristics; identity and status from the medicines register and assessment-report page	1 August 2026
Zuranolone, India	No CDSCO marketing approval record was located. Absent from all 66 published approval lists; the three documents a CDSCO-domain search returned were text-checked and none names the drug; no indexed CTRI record	Would require a CDSCO or DCGI new-drug permission, an approved label or an official list entry	1 August 2026
Brexanolone, United States, marketing status Discontinued	Indicated for the treatment of postpartum depression (PPD) in patients 15 years and older [see Clinical Studies (14)]; date of approval 17 June 2019; marketing status Discontinued	Drugs@FDA, NDA211371; indication from the single-column prescribing information, date from the approval letter	2 August 2026
Brexanolone (United States marketing status Discontinued), European Union	No European Union centralised procedure of any kind is on record for this active substance; the register was scanned whole and returns no row under any name. No European marketing status is on record	Would require an EMA assessment report or a Commission decision	1 August 2026

AGENT AND JURISDICTION	WHAT THE RECORD SAYS	AUTHORITY RELIED ON	STATUS TRUE AS OF
Brexanolone (United States marketing status Discontinued), India	No CDSCO marketing approval record was located. Absent from all 66 published approval lists; the single CDSCO-domain hit, expert-committee minutes dated 18 July 2024, is readable and does not name the substance; no indexed CTRI record. No Indian marketing status is on record	Would require a CDSCO or DCGI new-drug permission, an approved label or an official list entry	1 August 2026

Two properties must be stated in words or a reader will over-read the table. The register covers the European Union centralised procedure only, so a medicine authorised nationally through the decentralised or mutual-recognition route would not appear in it. What the empty row establishes is that no central application was ever concluded, by authorisation, refusal, withdrawal, lapse, revocation or suspension, since the register carries all those outcomes. That is narrower than "not available anywhere in Europe". For this substance the scope limit bites hard, and the confirming second instrument never ran on any gap row.

India second. The published CDSCO approval-list corpus is demonstrably incomplete, and no positive control establishes its sensitivity in the period these agents belong to: the one drug in this chapter whose Indian approval is known from outside the corpus, through expert-committee minutes, appears nowhere in the 66 documents, and the list declaring full-year coverage for that year carries entries dated after the approval it omits. A corpus that misses a drug the same regulator says it approved cannot prove another drug was not approved. So the sentence a candidate may write is not that these agents are unapproved in India, but that **no CDSCO approval record was located** as of August 2026, naming what was searched, on the authority of the Central Drugs Standard Control Organisation and the Drugs Controller General of India.

IN INDIA

Four things are mistaken for an Indian approval record. A CTRI registration records that a trial was registered, not that a drug was approved. Expert-committee minutes are advisory proceedings. A brand in a pharmacy list or a distributor catalogue is a commercial artefact. And an approval in Washington or Amsterdam is a fact about Washington or Amsterdam. The record statement: as of August 2026 no CDSCO approval record was located for either agent. This chapter holds no instrument that measured Indian availability either way.

- Mechanism: a neuroactive steroid modulates synaptic and extrasynaptic GABA-A receptors, reaching a tonic inhibition a benzodiazepine does not.

- A course, not a regimen: fourteen days orally, sixty hours by infusion, no second course described for either.
- A major depressive episode with an obstetric onset window, which is why the licence reads postpartum depression and the evidence goes no further.
- A status sentence names a jurisdiction, an exact indication, an authority and a date, or it is not written.

MNEMONIC

Subunit, Course, Clock, Border

Subunit separates this class from the benzodiazepines. Course is the finite exposure that follows. Clock is the onset window making a major depressive episode a postpartum one. Border is the rule that a regulator's record speaks only for its own territory.

The mechanism explains why the treatment is short, the onset window explains why the licence is narrow, and the jurisdiction explains why a true sentence about this drug can still be the wrong answer.

4 · Rapid-acting antidepressants: mechanism class and clinical positioning

A rapid-acting antidepressant is defined by a claim about time and licensed by a sentence about an indication. The class lives between those two objects. Its members share a response measured in hours to days rather than weeks, produced outside monoamine reuptake, and nothing else: not the receptor, the setting, the population each licence names, how long the treatment runs, or whether any regulator has concluded anything. Marks go in three places. Candidates write **rapid-acting antidepressant** as though a regulator had licensed the category; they collapse five members and three jurisdictions into one verdict; and they read the silence of one search as a fact about five products. The clinical accounts belong elsewhere: ketamine and esketamine in treatment-resistant depression to the Psychopharmacology II: emergencies, resistance, monitoring and interactions notes, their mood-disorder positioning to the Mood disorders: pharmacotherapy notes.

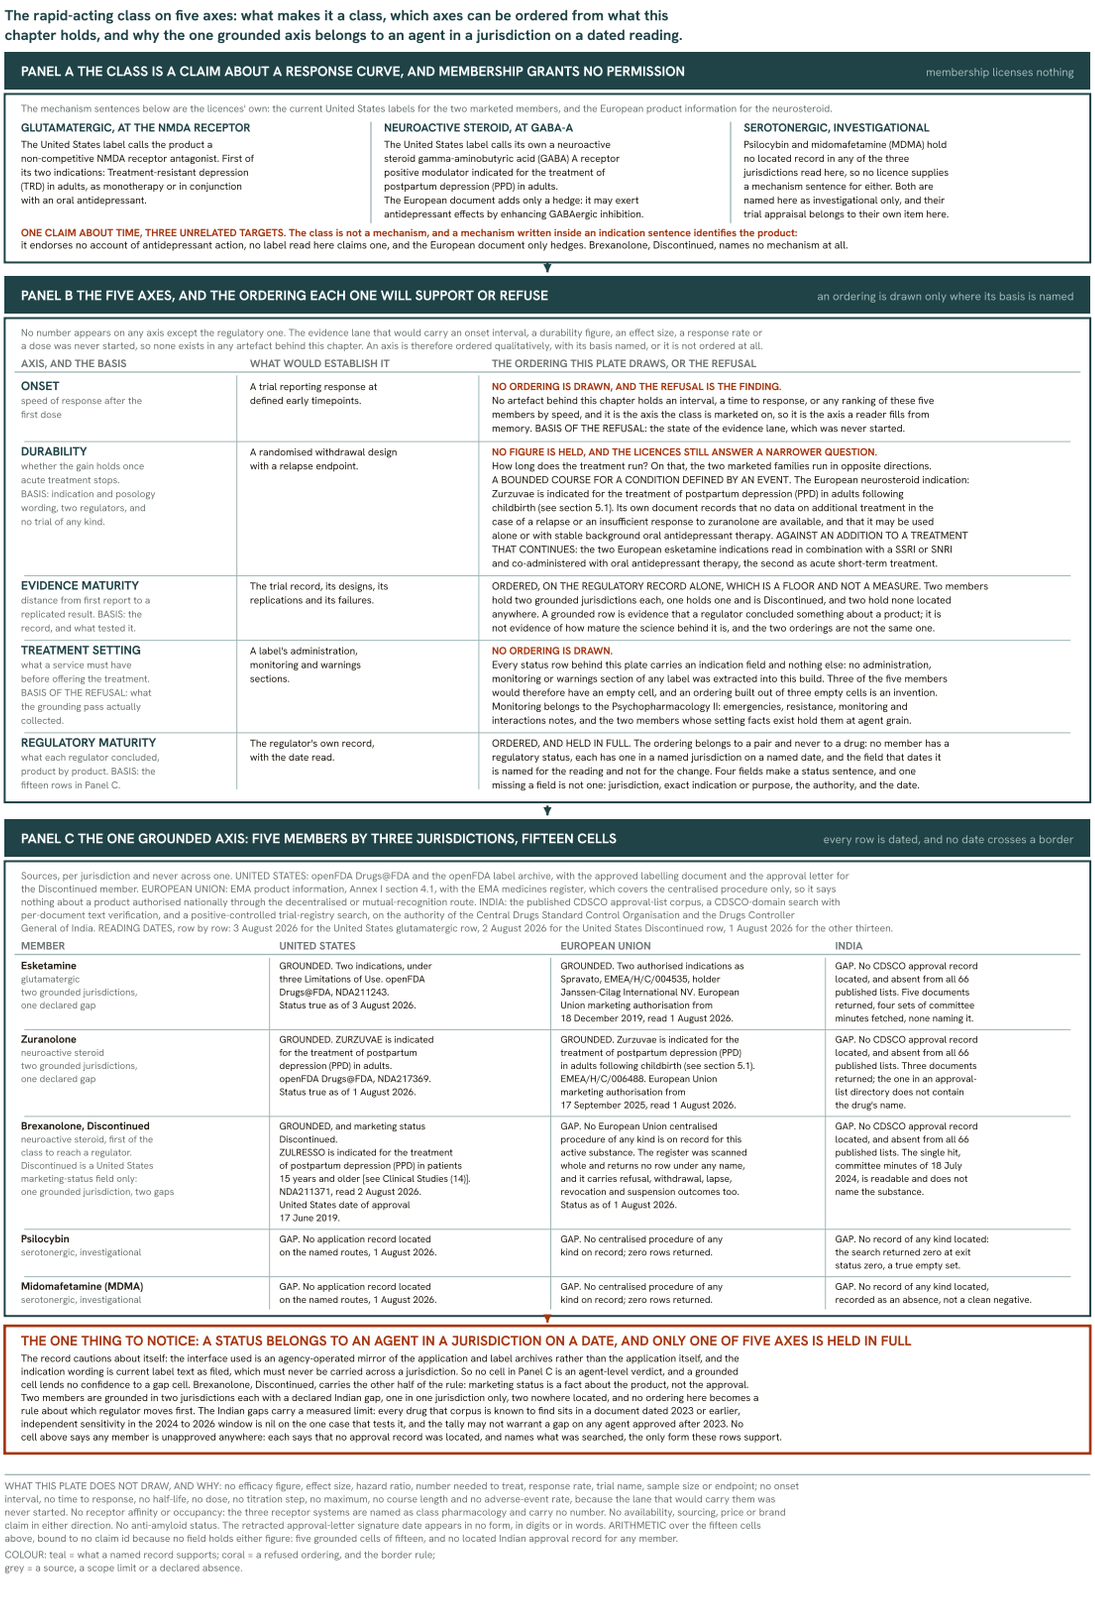


Fig 38.3 The rapid-acting class on five axes: only regulatory maturity is held in full, three of five not at all. Onset and treatment setting are refused orderings, each basis named. Durability is licence wording alone: a bounded course for a condition defined by an event, in adults following childbirth (see section 5.1), against an addition to continuing treatment, in combination with a SSRI or SNRI. Regulatory maturity belongs to an agent in a jurisdiction on a dated reading: five grounded cells of fifteen, with 17 June 2019 the United States date of approval for brexanolone, marketing status Discontinued. (openFDA Drugs@FDA, label archive and approval letter; EMA Annex I section 4.1, centralised procedure only; CDSCO approval lists and domain search, on CDSCO and DCGI authority.)

4.1 What makes this a class, and what the class promises

No regulator in these records authorises a rapid-acting antidepressant. Every authorisation names a product, an indication and a population, so the class is a construct built for a clinical purpose and carries no permissions. A sentence travelling from "this agent is rapid-acting" to where it may be used, in whom, or on whose authority has become a permission without passing through a document.

The promise sets the two questions the class is judged on. Conventional practice tolerates a latency of weeks and builds the service around that wait. Shorten it and two questions open: what happens to the gain when acute treatment ends, and what was the shortened wait supposed to buy?

One regulator has answered the second in writing. Among the **Limitations of Use** on the current United States label for esketamine stands the statement that The effectiveness of SPRAVATO in preventing suicide or in reducing suicidal ideation or behavior has not been demonstrated. The same label indicates the product for Depressive symptoms in adults with major depressive disorder (MDD) with acute suicidal ideation or behavior in conjunction with an oral antidepressant. The licence covers a population defined by suicidal ideation while declining any effect on the event that population is at risk of. Symptom change and prevention of an outcome are separate endpoints, and the regulator separated them. An answer arguing from speed of response to a suicide-prevention benefit makes the leap the licence refuses.

Notice what kind of negative that is, because a candidate trained to distrust categorical negatives may mistake it for one. It states what a trial programme demonstrated, printed inside a licence by the authority that granted it, so it is a positive fact about an evidence base; a claim that a drug is unapproved somewhere needs a regulator's affirmative statement and appears nowhere here without one. The same label's SPRAVATO is not approved as an anesthetic agent is again a licence describing its own boundary.

RULE **Rapid-acting describes a response curve and never grants a permission.** Every sentence about what may be given, to whom, where and on whose authority must be sourced to one named regulator's record for one named product, with the date read. The class name cannot stand in for a document.

4.2 Three mechanistic families, and the mechanism sentence a licence gives you

Three receptor systems sit under one heading: the glutamatergic agents at the NMDA receptor complex, the neuroactive steroids as positive modulators at the GABA-A receptor, and the serotonergic psychedelics, investigational throughout, at serotonin 2A receptors. One claim about time, three unrelated targets: this class is not a mechanism, and an answer offering one has answered another question.

Two of the three have a mechanism written into the licence. Esketamine's United States label calls the product **a non-competitive NMDA receptor antagonist**, spelling the receptor name out in full, indicated for Treatment-resistant depression (TRD) in adults, as monotherapy or in conjunction with an oral antidepressant. Zuranolone's United States label calls its own **a**

neuroactive steroid gamma-aminobutyric acid (GABA) A receptor positive modulator indicated for the treatment of postpartum depression (PPD) in adults. A mechanism inside an indication sentence identifies the product; it endorses no account of antidepressant action, and no label read here claims one.

The third teaches it from the other side. Brexanolone, whose United States **marketing status** is Discontinued, carries an indication naming no mechanism: ZULRESSO is indicated for the treatment of postpartum depression (PPD) in patients 15 years and older. The neuroactive-steroid account of it comes from the literature, not the licence. Asked what a licence says, answer only that.

No receptor affinity, occupancy figure, onset interval, half-life or comparison of potency or speed between these families is held by any artefact behind this chapter, because the evidence lane that would carry them was never completed. Writing one from memory would be a fabricated citation, the one defect no later pass repairs.

4.3 The five axes, and which of them anything here fills

A class map is useful only if it is honest about its columns, and here one of five axes is held in full, one narrowly, and three not at all. A table hiding that invites a reader to fill the empty columns from memory and quote the result as though the map supplied it.

AXIS	WHAT IT WOULD MEASURE	WHAT WOULD ESTABLISH IT	WHAT THESE NOTES HOLD
Onset	Speed of response after the first dose	A trial reporting response at defined early timepoints	Nothing: no interval, no time to response, no ranking by speed
Durability	Whether the gain holds once acute treatment stops	A randomised withdrawal design with a relapse endpoint	No figure. Only licence wording, separating a bounded course from a continuing treatment
Evidence maturity	Distance from first report to replicated result	The trial record, its designs, replications and failures	Only the regulatory record and the instruments that tested it: a floor, not a measure
Treatment setting	What a service must have before offering the treatment	A label's administration, monitoring and warnings sections	Nothing directly; never extracted. Monitoring belongs to the chapters named above
Regulatory maturity	What each regulator concluded, product by product	The regulator's own record, with the date read	Held in full, at the grain of one agent in one jurisdiction

The columns a candidate most wants to fill are the two the class is marketed on, so the defensible answer runs the map backwards: from what each regulator concluded, to what the licence implies about treatment length, then naming the rest unestablished.

4.4 Durability: the axis two licences answer in opposite ways

Nothing here quantifies durability, and the licences still answer the question. They answer the narrower version a service plans around: not how long a response lasts, but how long the treatment runs. On that the two marketed families run in opposite directions in the indication text.

The neurosteroid licences are built around a bounded situation. The European authorisation reads Zurzuvae is indicated for the treatment of postpartum depression (PPD) in adults following childbirth (see section 5.1). The United States indication for the same molecule reads ZURZUVAE is indicated for the treatment of postpartum depression (PPD) in adults. A condition defined by its relation to an event implies a treatment with an end, and the European text names the event where the American does not.

The glutamatergic member's European licence runs the other way, indicating esketamine in combination with a SSRI or SNRI for treatment-resistant major depressive disorder, and in its second authorised indication co-administered with oral antidepressant therapy for a moderate to severe episode as acute short-term treatment, for the rapid reduction of depressive symptoms, which according to clinical judgement constitute a psychiatric emergency. Two facts sit there: both European indications license the agent inside a continuing antidepressant treatment rather than instead of one, and the second carries its own durability limit in the words acute short-term treatment.

The class therefore splits along a line unrelated to receptor: a bounded course for a bounded condition, against an addition to a treatment that continues. Both say how long a prescription runs, not how long a response lasts, which decides staffing and review intervals. The United States monotherapy line belongs to esketamine's own section here.

One demographic comparison is available. Brexanolone, marketing status Discontinued, is licensed, in the United States, in patients 15 years and older, while the indications quoted above for both marketed members name adults. The class's only reach below adulthood sits on the member that is Discontinued, a fact about a licence rather than an option to offer anyone.

> **GOOD TO REMEMBER**
>
> Two questions answer the planning problem from the licence alone, without a trial: is this agent licensed as an addition to something that continues, and is the condition it names defined by an event that ends? Different clinical commitments, and the patient should be told which is starting.

4.5 Regulatory maturity belongs to a pair, not a drug

Four fields make a status sentence and one missing a field is not one: the jurisdiction, the exact indication or purpose in the source's words, the authority the claim rests on, and the date it was true. That last field's name is worth using, **status true as of**, because it dates the reading and not the change. The United States rows here were read on 3 August 2026 for esketamine, 1 August for zuranolone and 2 August for brexanolone, marketing status Discontinued, and every European and Indian row on 1 August 2026.

Five members and three jurisdictions make fifteen cells, and they disagree inside one agent: no member has a regulatory status, each has one in a named jurisdiction on a named date. The record says so of itself: the interface used is an agency-operated mirror of the application and label archives rather than the application itself, and the indication wording is current label text as filed, which must never be carried across a jurisdiction.

MEMBER	GROUNDED RECORDS	DECLARED GAPS	WHAT MAY NOT BE WRITTEN
Esketamine, glutamatergic	United States, two indications under three Limitations of Use, NDA211243, read 3 August 2026. European Union, two indications as Spravato, EMEA/H/C/004535, holder Janssen-Cilag International NV, from **18 December 2019**	India: no CDSCO approval record located	Either indication carried across, and any Indian status either way
Zuranolone, neurosteroid	United States, postpartum depression in adults, NDA217369. European Union, the same with following childbirth added, as Zurzuvae, EMEA/H/C/006488, holder Biogen Netherlands B.V., from **17 September 2025**	India: no CDSCO approval record located	The European wording quoted as the American indication, and the register's decision column read as an authorisation date
Brexanolone, neurosteroid, marketing status Discontinued	United States only, postpartum depression in patients 15 years and older, NDA211371, date of approval **17 June 2019**, read 2 August 2026	European Union: no centralised procedure on record; India: no CDSCO record located	That its approval has lapsed, and any inference from Discontinued about the gap rows
Psilocybin, investigational	None located	United States, European Union and India, all three declared	Any categorical negative, and any upgrade of investigational status to an authorisation
Midomafetamine (MDMA), investigational	None located	United States, European Union and India, all three declared	The same two, plus transfer from the other investigational member

The investigational pair belongs to the psychedelic-assisted therapy item in this chapter, which owns their trial appraisal. They appear only because a class comparison that drops its immature members flatters the class.

TRAP **Collapsing the map to one verdict per agent.** Two members hold two grounded jurisdictions each with a declared Indian gap, one holds one grounded jurisdiction and two gaps, two hold none anywhere. "Esketamine is approved" is true of two records, unsupported for a third, and unmarkable because it names no regulator. Never write an approval verb without the jurisdiction, source and date beside it, and never let a grounded row lend confidence to a gap row.

4.6 Brexanolone, Discontinued: a probe miss, a date set by scheduling, and an empty key

One row has been wrong in every interesting way and repaired each time. It also carries the rule for every sentence about it: teach this agent and carry its marketing status in the same breath, since marketing status is a fact about the product and not about the approval, and this product is Discontinued while its approval stands.

The row read GAP, on a verdict of no record found, until multi-route probing retrieved it. The verdict was a probe miss and not an absence: the record carries no agency-standardised name block, so queries through generic name, substance name and brand name returned not-found while the query reading the product's active-ingredient field returned it. The record's own sentence transfers: a decided action is not a licence to ignore a record that turns up.

Then the date, a controlled-substance lesson in clerical dress. The United States date of approval for brexanolone, marketing status Discontinued, is 17 June 2019, and the approval letter did not set it by its own signing. The letter states that the date of approval shall be the date on which the Drug Enforcement Administration publishes a notice in the Federal Register announcing the interim final scheduling of brexanolone, marketing status Discontinued. The original submission carries an earlier status date, that signature and not the approval; no number for it is printed here.

The third repair is worth the most and looks like housekeeping. This row's canonical earliest-approval key is deliberately absent, because writing the submission's signature date under a field name meaning first approval would put a retracted number back into the field every consumer reads. Rarer than getting that right, it writes down the cost. The carry check reads only the European, Indian and canonical earliest-approval keys, so it cannot see where this row's approval date now lives; the row sits outside that check until it learns the key, and the pass's claim to compare every date field is no longer true of this file. A correct repair left a control blind, and the record says so.

It is also the one indication here no instrument could test. All three text instruments returned unresolved extraction rather than a pass, for a documented reason: with no name block the label archive cannot be queried by application number, so the indication came from the approved labelling document instead; the pass's own discipline is that a failed extraction is unresolved extraction and never a pass. How it was read matters too. The indication was taken from the single-column prescribing information, not the two-column highlights block, whose extract interleaves the columns. Interleaved columns produce fluent sentences no document contains.

5 of 15	0 of 5	3 of 3	2 of 2
GROUNDED CELLS, FIVE MEMBERS BY THREE JURISDICTIONS	INDIAN APPROVAL RECORDS LOCATED FOR THESE MEMBERS	BREXANOLONE, DISCONTINUED, US TEXT CELLS UNTESTED, NOT PASSED	EU ESKETAMINE INDICATIONS CARRIED AGAINST SOURCE

The fourth figure counterweights the third: The completeness instrument found two indications in the European source document and two on the row, and the instrument set was validated by planting the real historical defect on that row, a dropped second indication, and confirming the guard fired. A pass from an instrument tested against the failure it exists to catch is worth something.

TRAP **Reading an untested cell as a clean one.** Three cells here were never tested, because the source text could not be retrieved by the route the instrument uses, and are recorded as unresolved rather than as passes. A reader who scans a matrix for the word fail, finds none and calls the column sound has learned nothing. The mirror error costs as much: an unresolved text check does not undermine a status resting on the labelling document and the approval letter.

4.7 The Indian column, and what nine search hits established

Every member has the same Indian entry: no located CDSCO marketing approval record. At class level that is one instrument's result five times over, so all five cells carry that one search's sensitivity, which the record measures rather than assumes.

What was opened is the lesson. For esketamine, a domain search returned five documents, four sets of subject expert committee minutes for neurology and psychiatry and one international cell file; all four were fetched, all carry an extractable text layer, and none names the drug, so all were search-index false positives and not evidence. For zuranolone, three came back, and the one in an approval-list directory was downloaded and its text does not contain the drug's name, as with both sets of minutes. For brexanolone, marketing status Discontinued, the single hit, committee minutes dated 18 July 2024, was fetched, is readable, and does not name the substance. Nine hits, eight documents opened, nine false positives, the fifth esketamine document scored one without being fetched: a search hit is not a document naming your drug.

The two investigational members returned nothing, and even that is graded. The domain search returned zero results at exit status zero, a genuine empty result set rather than a failed query, and the finding is still recorded as the absence of any located record rather than a clean negative, which would overstate what an instrument with a measured miss rate and three failed positive controls delivers.

Two sensitivity facts belong in any Indian sentence about a recent agent, and the first has a name, **era mismatch**. Every drug this corpus is known to find sits in a document dated 2023 or earlier, and independent sensitivity in the 2024 to 2026 window is nil on the one case testing it; the positive-control tally may not warrant a gap on any agent approved after 2023.

The authority is named, and its substitutes with it. The claim would rest on the Central Drugs Standard Control Organisation and the Drugs Controller General of India, and committee minutes, trial no-objection certificates, international cell confirmations, CTRI registration, brand

listings, pharmacy availability and manufacturer copy are none of them marketing authorisation. A CTRI entry records that a trial exists, not permission to market. The rule closing the column sits in each of these rows: no India row here asserts that a drug is not approved in India; each asserts that no CDSCO approval record was located, and names what was searched.

IN INDIA

Write the Indian column in this order. No CDSCO marketing approval record was located for any member as of 1 August 2026, on the three instruments named above, and the corpus has never been shown to find an approval from 2024 onward. Do not write that any is not approved in India, which is stronger than the evidence supports, and do not write that any is approved, because nothing retrieved says so. For a family: there is no located approval record to point to, and licensing is a question for the regulator's record, not a brand listing.

4.8 Choosing across the class, and what a complete answer contains

Five statements answer it in whatever form it is asked.

- The class is a claim about a response curve, not a licensed category, so membership licenses nothing.
- Three unrelated receptor systems share one heading, and a licence naming a mechanism identifies a product rather than endorsing an account of action.
- Of five axes a class map needs, only regulatory maturity is held in full, durability only in the narrow form licence wording settles.
- Regulatory maturity belongs to an agent and a jurisdiction on a stated date; a grounded row never lends confidence to a gap row.
- The Indian column is five copies of one search result with a measured sensitivity: a statement about a search, never about a product.

MNEMONIC

Curve, Country, Course, Corroboration

Curve is what the class name describes. Country is the jurisdiction every status sentence must name and never leave. Course is durability in the form a licence answers, how long the treatment runs. Corroboration is the last question before quoting a cell: was it tested, tested against the failure it exists to catch, or merely unmarked?

The class name describes a curve and grants nothing; status belongs to an agent and a jurisdiction on a dated reading; and an untested cell is not a clean one.

5 · Dual orexin receptor antagonists (DORAs) and sodium oxybate in sleep psychiatry

Four agents that one sleep service may use in the same week hold, between them, no single shared pattern of national approval. Three **dual orexin receptor antagonists** and one oxybate show that regulatory status belongs to an agent inside a jurisdiction and never to a class. The

pharmacology is not rebuilt here: the orexin mechanism, the agents, dosing, adverse effects and insomnia prescribing belong to the Dissociative, conversion and anxiolytic/hypnotic pharmacology notes, and narcolepsy, cataplexy, orexin biology and the disease-specific use of sodium oxybate belong to the Eating and sleep-wake disorders notes. What is dated and new is the regulatory record. Three errors cost most: one class verdict for four agents, an approval carried across a border, and a sentence read as complete because it ends in a full stop.

5.1 Four fields, and why this group needs them most

A status sentence is checkable only if it carries the jurisdiction, the indication in the regulator's own words, the named source, and the date the reading was true. The rule that follows governs the chapter: a row may be written only in the jurisdiction it was grounded in, and none crosses a boundary. A declared absence is written as an absence, naming what was searched and whose authority a claim would rest on; an inferred figure is not written at all.

At the grain of one agent in one country these four generate twelve records that disagree. Seven carry an affirmative located record and five are declared absences, and which is which follows neither the agent, the class, the year nor the indication; the tally is computed across these rows. The three orexin antagonists are licensed, where they are, for insomnia, and the oxybate for narcolepsy with cataplexy or the excessive daytime sleepiness of narcolepsy.

RULE **An insomnia approval is never transferred to narcolepsy or cataplexy, and a narcolepsy approval is never transferred to insomnia.** Both families act on the same sleep-wake machinery and are prescribed by the same clinician, which is why the licences get swapped in an answer. Quote the indication belonging to that product in that country.

5.2 The United States: four records, four dates, one reading date that is not the others

All four carry an affirmative American record, read from the openFDA Drugs@FDA interface with a **structured product label** named by identifier and effective date, current for three rows and not for the oxybate. Three resolve to one application; the fourth does not, which is 5.3.

AGENT	APPLICATION AND LABEL	INDICATION AS FILED, QUOTED	EARLIEST APPROVAL-STATUS DATE IN THE RECORD
Suvorexant	NDA204569, label 21306af5 effective 10 March 2025	BELSOMRA (suvorexant) is indicated for the treatment of insomnia characterized by difficulties with sleep onset and/or sleep maintenance.	13 August 2014
Lemborexant	NDA212028, label c82838dc effective 24 February 2025	DAYVIGO is indicated for the treatment of adult patients with insomnia, characterized by difficulties with sleep onset and/or sleep maintenance	20 December 2019
Daridorexant	NDA214985, label 10e24e38 effective 26 February 2026	QUVIVIQ is indicated for the treatment of adult patients with insomnia, characterized by difficulties with sleep onset and/or sleep maintenance	7 January 2022
Sodium oxybate	Four applications, two original; text re-read from label 3839d447 effective 1 July 2025, whose own application number is NDA021196, not the row's current label	Xyrem is indicated for the treatment of cataplexy or excessive daytime sleepiness (EDS) in patients 7 years of age and older with narcolepsy.	**17 July 2002**

Three differences are examinable. The three insomnia texts differ: two name adult patients and the earliest names no population, a difference between three documents rather than an inference about who may be treated. The oxybate row does not call itself an **orexin receptor antagonist** as the other three do; its label calls it a central nervous system depressant. And the last column is named for the field it came from: whether any is a first marketing approval these records do not say.

The record cautions about itself: openFDA is an agency-operated mirror of Drugs@FDA and of the label archive rather than the application itself, and the indication wording is current label text as filed, which is not a statement about any other jurisdiction and must never be carried across one. The reading dates differ: the three orexin rows carry a status true as of 1 August 2026 while the oxybate row carries 3 August 2026, because it was repaired in between. Harmonising the newer date onto the others asserts readings nobody made.

5.3 Right citation, right date, wrong product

The oxybate row is the most instructive failure here, because nothing about it looked wrong. The cited application and the date were both correct and both belonged to Xyrem. The indication text was not: it was the text of LUMRYZ, a different application by a different sponsor, approved in 2023, whose label reads almost identically in structure at 349 characters against 351 and differs mainly in the product name. Comparing lengths would have passed it; checking the retrieved document's own application number against the one cited caught it.

The vulnerability is structural: three of these four return one application, so there is nothing to blend; sodium oxybate returns four. Adjudication had already established that Xyrem, Xywav and Lumryz are three separate applications and that Xywav is not even a single-entity sodium oxybate product; that finding could not reach the ledger because its regeneration was parked, so the defect stood for a day after being solved elsewhere. The row is checkable now because it records an identity test anyone can repeat.

5.4 The European Union: two authorised, two absences, and a cut that closed with a full stop

Two agents carry a European centralised authorisation and two do not. For the two that do, indication text comes from Annex I of the **Summary of Product Characteristics**, section 4.1 Therapeutic indications, read from the cached product-information document; the register and assessment-report page corroborate identity and status only. Daridorexant: QUVIVIQ is indicated for the treatment of adult patients with insomnia characterised by symptoms present for at least 3 months and considerable impact on daytime functioning. The oxybate: Treatment of narcolepsy with cataplexy in adult patients, adolescents and children from the age of 7 years. Both are Authorised, daridorexant under EMEA/H/C/005634 and the oxybate under EMEA/H/C/000593, each with a **marketing authorisation holder** the register names by company and not by country: Idorsia Pharmaceuticals Deutschland GmbH and UCB Pharma Ltd. Both dates are of a declared kind, the initial European marketing authorisation as a Commission decision: **29 April 2022** and 13 October 2005, each file carrying a later Commission decision in 2026 and 2025.

> **GOOD TO REMEMBER**
>
> Whether a patient must have cataplexy to fall inside the licence depends on which regulator you are reading: one text reaches excessive daytime sleepiness in narcolepsy without requiring cataplexy, the other names narcolepsy with cataplexy. When a colleague quotes an oxybate indication, the question is not what it says but which country's document it came from.

That row carries this section's instrument lesson. Until it was corrected that row read "Treatment of narcolepsy with cataplexy in adult patients." The words "adolescents and children from the age of 7 years" had been dropped and the surviving prefix closed with a full stop the source does not have. The same correction fixed a second defect: the indication had come from the published register while the assessment-report page was cited as an independent second instrument, though the register is generated from those pages, so the corroboration was circular and blind to the stale summary layer that had failed.

TRAP **Treating a sentence that ends in a full stop as one that ends where the source ends.** A truncation landing mid-word announces itself; one landing at a clause boundary and closing with valid punctuation is undetectable by eye, and it removed a paediatric population from an authorised indication in this ledger. The instrument that catches it asks one question: is the stored text a prefix of the primary text? Learn the question, not the incident.

For the two with no European record the absence is written as an absence, its strength stated. The agency's own medicines register, generated on 1 August 2026, was scanned whole across all 39 columns of all 2730 medicine rows, and each substance returns zero rows under any of its names. The register carries Authorised, Refused, Application withdrawn, Withdrawn, Lapsed, Expired, Revoked, Suspended and Opinion outcomes, so absence means no central application of any kind was ever concluded, including a refused or abandoned one; what would close the row is an agency assessment report or a Commission implementing decision. The authority a claim would rest on is named, and "approved internationally" is not a permitted substitute for a named regulator.

Then the limit: the register covers the European **centralised procedure** only, and a product authorised nationally through the decentralised or mutual-recognition route would never appear in it. That limit is printed identically on every European absence here but does not bite equally, and on both these agents it is graded as biting weakly, because a small molecule outside the mandatory centralised scope could hold a national authorisation this instrument cannot see. Two weaknesses travel with the finding: the confirming second instrument ran on all eight grounded rows and none of the eight absences, so no European absence here has independent corroboration, and the assessment-report page would anyway be a second rendering of the register's own source.

5.5 The ladder that does not exist

Lemborexant carries a located Indian approval record and a declared European absence. That defeats the assumption almost every candidate brings to a recent-advances question: that maturity runs America, then Europe, then India. Here the middle rung is the empty one, and these four agents run three patterns across three jurisdictions.

The Indian record is affirmative. The approval-list cell reads, whole, Treatment for insomnia. Lemborexant tablets 5mg/10mg. It sits at serial number 12 of a 2021 approval list, one document in a corpus of 66 distinct published approval documents, deduplicated by content hash, discovered by four mapping passes and downloaded directly, all 66 yielding an extractable text layer, so no absence in this corpus is an artefact of an unreadable file. The strengths there are a formulation fact from the regulator's own name column, quoted because the discipline is to carry a cell whole; they are not a dosing recommendation, and no dose or maximum appears here.

The date needs the same care. The value is **1 July 2021**, and the field says what is printed: the figure in the column headed **Date of issue**. An earlier version claimed to quote a label the document does not carry, and was narrowed. Corroboration is thinner than it looks: two documents are cited, but they are one running list in two renderings rather than two independent documents, and the second lies outside the 66-document corpus, found by a drug-specific search

and appearing in no mapping pass, so the corpus discovery step was incomplete though the fetching step was not.

TRAP **Answering at the grain of the agent instead of the agent and the country together.** "Lemborexant is approved" is not a shorter version of a true sentence but a different and unverifiable one. Treat each of these twelve records as its own row, and answer any availability question with the jurisdiction, the record and the date. The mirror-image error costs as much: inferring a European or American absence from an Indian gap, when here the Indian record is the affirmative one.

5.6 The three Indian absences, and the one built on documents that exist

The other three carry no located Indian marketing approval, and the absences are not equal. Suvorexant's is strongest, resting on documents that exist rather than on documents that do not. The row records the file at pre-approval stage: CDSCO subject expert committee minutes for neurology and psychiatry of 24 September 2025 record a named manufacturer proposing permission to manufacture and market suvorexant tablets, the committee recommending permission to conduct a bioequivalence study, and minutes of 26 November 2025 record a second manufacturer with the same proposal, the committee recommending a Phase III clinical trial. Both are applications, and an application is not an approval.

That row's lead sentence once read that the agent was not approved in India at retrieval, and it was withdrawn: pre-approval activity by two named firms is evidence that those two firms had not yet been granted permission, not proof that no other holder exists, in a corpus this build itself shows to be incomplete. The evidence did not change; only the claim it was asked to carry was narrowed, and the row remains a declared absence.

The other two rest on searches that returned nothing usable, each naming what it searched. For daridorexant the substance is absent from all 66 published lists; four committee minutes returned by a domain search were all fetched, are all readable and none names the drug, and a fifth hit in an approval-list directory does not contain the string either, so all five were search-index false positives; no indexed CTRI trial record was located. For the oxybate the five domain hits are all International Cell written confirmation files naming Indian manufacturers, export-related documents that are not Indian marketing authorisation; no committee minute, no clinical-trial approval and no indexed CTRI record was located. A **written confirmation** is likeliest to be mistaken for an approval, being a real regulatory document about a real manufacturer that says nothing about the Indian market.

The authority that would settle any of these is the Central Drugs Standard Control Organisation and the Drugs Controller General of India, and committee minutes, clinical-trial no-objection certificates, International Cell written confirmations, CTRI registration, brand listings, pharmacy availability and manufacturer copy are none of them marketing authorisation. A **CTRI registration** records that a trial exists, not permission to market. Two sensitivity facts then bind any Indian sentence here. Every one of the seventeen positive controls this corpus finds sits in a document dated 2023 or earlier, and measured independent sensitivity in the 2024 to 2026 window is 0 of 1, so the corpus has never been shown to find a recent approval; and the positive-

control tally of 17 of 20 may not be cited as warrant for an absence on any agent approved after 2023. Absence from this corpus is evidence and not proof, and no Indian row in this build asserts that a drug is not approved in India: each asserts that no CDSCO approval record was located, and names what was searched.

IN INDIA

Write it one agent at a time. Lemborexant: a CDSCO approval-list entry for insomnia at serial number 12 of a 2021 list, date of issue 1 July 2021. Suvorexant: no CDSCO marketing approval located, the file documented at pre-approval stage on 2025 committee minutes; an application is not an approval. Daridorexant and sodium oxybate: no CDSCO approval record located, and what was searched is named. Do not write that any of them is not approved in India, and do not write that any is approved without naming the located record.

5.7 What certified these rows, and what a null cell is worth

These twelve records were tested rather than trusted, on three instruments: is the stored text verbatim against the primary document, is it that text cut short, does the count of indications carried match the document's? That is 36 cells, of which 21 are applicable and all 21 pass, with no failures and nothing unresolved.

1 of 2730	21 of 21	349 of 349	1 of 66
EU REGISTER MEDICINES NAMING OXYBATE	APPLICABLE FALSIFICATION CELLS PASSING	CHARACTERS, US OXYBATE ROW VERSUS LABEL	CDSCO LISTS WITH A LOCATED ENTRY

The fifteen inapplicable cells matter as much, and the reason for skipping them defines a declared gap: a row that is a declared absence carries no indication text, so an instrument comparing text against a primary document has nothing to test, because a gap is not a claim about a product but a claim about a search. An absence can be certified only as honest about its own reach.

Two instruments were proved on defects taken from this group, not invented ones. The prefix instrument was tested by replanting the real European oxybate truncation and returned a failure; the verbatim instrument by replanting the real Indian misattribution, in which a flattened text extraction attached the obesity indication belonging to the next serial number, cetilistat, to the lemborexant row, and it also returned a failure. That defect was repaired by reading the raw column layout of the source table instead of the flattened text line. An instrument never shown to fire on a known defect is not evidence.

The Indian cell is certified on a geometric reader's agreement with the drawn rules, its locator recorded to document, page, serial number and column coordinates, while the second document that row cites is absent from the cached corpus altogether. Certification concerns the reading, corroboration the sources.

5.8 The four records side by side, and what a complete answer contains

Reproduce this grid and the item is answerable in any form.

AGENT	UNITED STATES	EUROPEAN UNION	INDIA
Suvorexant	Located label indication for insomnia, NDA204569	No centralised procedure of any kind on record; scope limit weak	No approval located; file at pre-approval stage on 2025 minutes
Lemborexant	Located label indication for adult insomnia, NDA212028	No centralised procedure of any kind on record; scope limit weak	Approval-list entry for insomnia, serial 12, date of issue 1 July 2021
Daridorexant	Located label indication for adult insomnia, NDA214985	Authorised, EMEA/H/C/005634; indication carries a duration and a functioning criterion	No approval located; four minutes read, none names the drug
Sodium oxybate	Located indication: cataplexy or excessive daytime sleepiness in narcolepsy from age seven	Authorised, EMEA/H/C/000593; indication is narcolepsy with cataplexy including children from age seven	No approval located; the five documents found are export written confirmations

Reduced to four sentences:

- Status belongs to an agent inside a jurisdiction on a date, so twelve records is the right number and one class verdict is wrong.
- The two families carry different diseases in their licences, and their indications are never substituted for one another.
- The commonest silent defect is a regulator's sentence cut at a clause boundary and closed with a full stop.
- A declared absence is a statement about a search, naming the instrument, its reach, and whose authority would make it a negative.

MNEMONIC

Agent, Border, Cell, Cut

Agent: one drug in one country, never a class. Border: no indication, date or verdict crosses a jurisdiction. Cell: quote the regulator's cell whole, including the population at its end. Cut: ask whether a quoted indication is the primary text closed off tidily.

The licence belongs to a drug in a country on a date, and a sentence that ends in a full stop has not thereby ended where its source did.

6 · Xanomeline-trospium and muscarinic agonist antipsychotics

One capsule holds a muscarinic agonist and a muscarinic antagonist, and that pairing is the design rather than an accident of formulation. Two disciplines are examined at once. The

pharmacological question is how a treatment for psychosis that gives up dopamine blockade is supposed to work, and why the second molecule is there to be kept out of the brain rather than let into it. The regulatory question is why three national records of the same drug read so differently. Marks are lost in three predictable places: candidates treat the antagonist as a central anticholinergic; they carry an American indication across a border; and they read the silence of an instrument never shown able to speak. Receptor principles for the muscarinic system belong to the Psychopharmacology I: principles and core drug classes notes and are not rebuilt here.

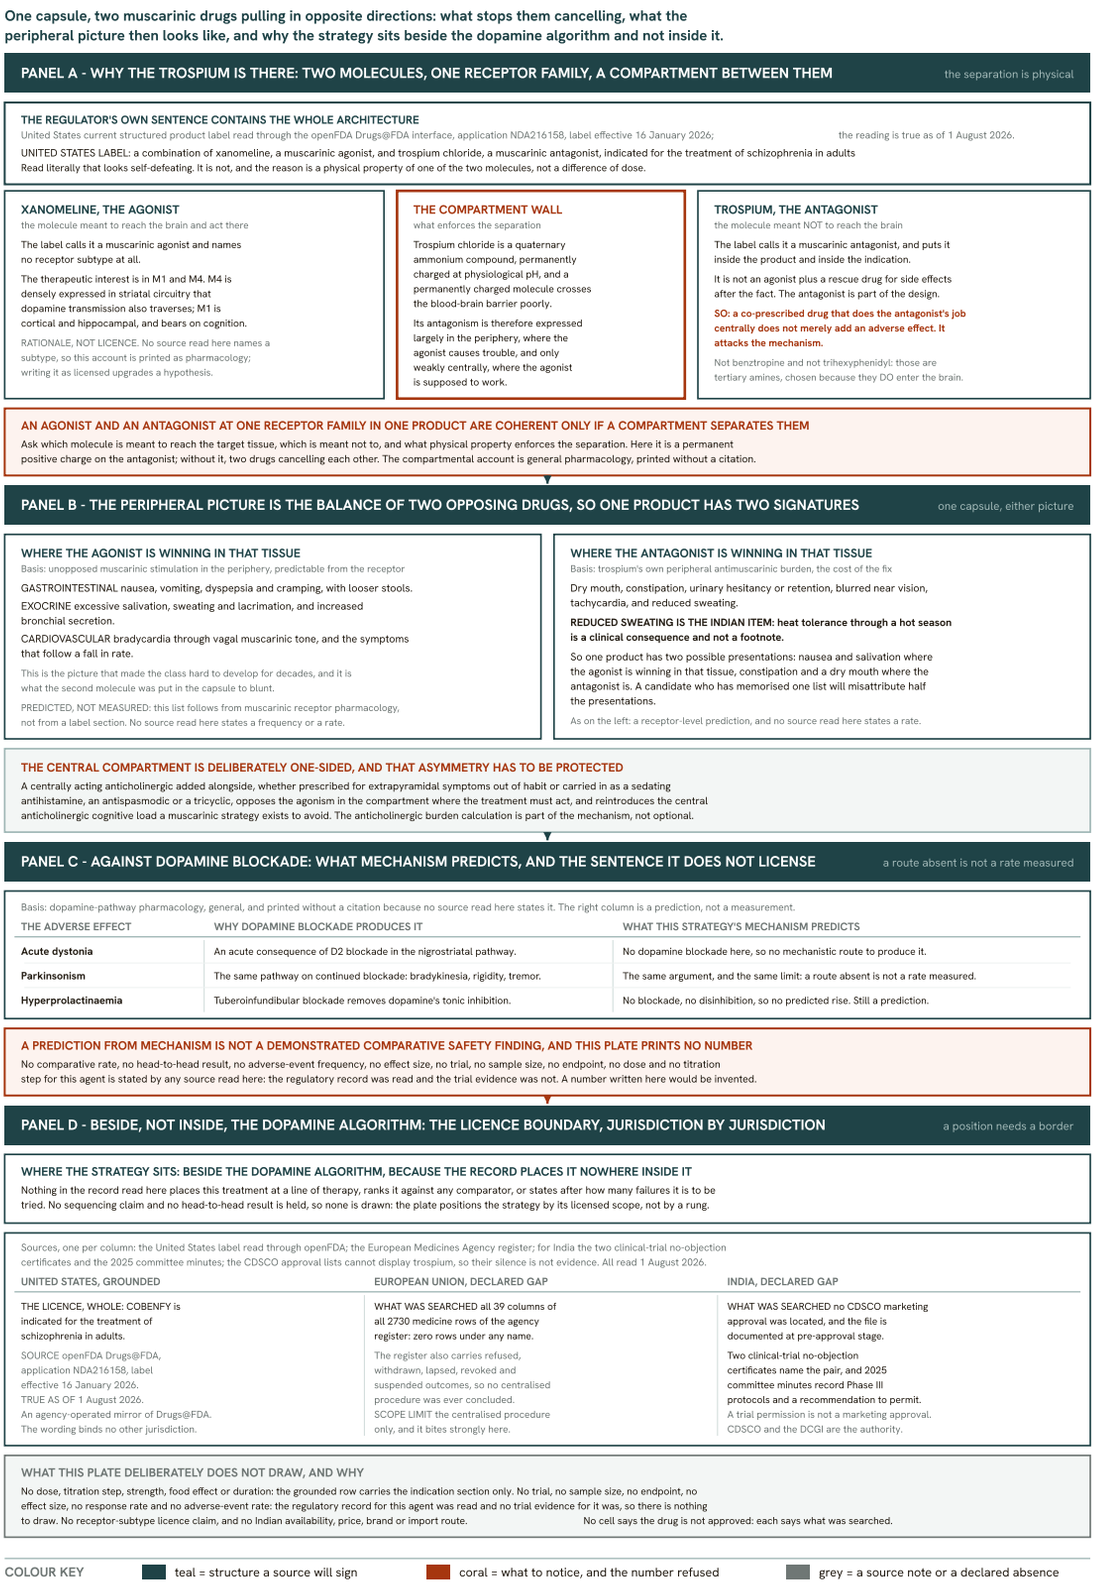


Fig 38.4 Muscarinic agonism with a peripherally restricted antagonist: why the trospium is in the capsule, and where the strategy sits. The label describes a combination of a muscarinic agonist and a muscarinic antagonist, and the pairing is coherent because trospium is permanently charged and stays peripheral, so one product can present with either signature, both predicted from the receptor and neither with a rate. The licence is the treatment of schizophrenia in adults in one jurisdiction; the European Union and India are drawn as declared gaps naming what was searched. (United States label via openFDA, NDA216158, effective 16 January 2026; European Medicines Agency medicines register; for India two trial no-objection certificates and 2025 committee minutes; read 1 August 2026.)

6.1 An agonist and an antagonist in one capsule, and why that is not a contradiction

Start with the sentence the regulator wrote; it contains the whole architecture. The product is described on its current United States label as a combination of xanomeline, a muscarinic agonist, and trospium chloride, a muscarinic antagonist, indicated for the treatment of schizophrenia in adults. Read literally that looks self-defeating; a candidate who says so has noticed the right thing and drawn the wrong conclusion.

The pairing is coherent because the two molecules are separated by a compartment and not by a dose. Xanomeline, the **muscarinic agonist**, is intended to reach the brain and act there. Trospium, the **muscarinic antagonist**, is a quaternary ammonium compound, permanently charged at physiological pH, and a permanently charged molecule crosses the blood-brain barrier poorly. Its antagonism is therefore expressed largely in the periphery, where the agonist causes trouble, and only weakly centrally, where the agonist is supposed to work.

Two things follow. The combination is not an agonist plus a rescue drug for side effects after the fact; the antagonist is inside the product and inside the indication. And a co-prescribed drug that does the antagonist's job centrally does not merely add an adverse effect. It attacks the mechanism.

RULE **An agonist and an antagonist at the same receptor family in one product are only coherent if a compartment separates them.** Ask which molecule is meant to reach the target tissue, which is meant not to, and what physical property enforces the separation. Here it is a permanent positive charge on the antagonist; without it, two drugs cancelling each other.

6.2 What the muscarinic strategy replaces, and where it sits

Every antipsychotic algorithm a postgraduate has learned is built on occupancy of the dopamine D2 receptor. This strategy is not on that axis; it engages the cholinergic muscarinic system instead, and the therapeutic interest sits with the M1 and M4 subtypes, M4 because it is densely expressed in striatal circuitry dopamine transmission also traverses, M1 because of its cortical and hippocampal distribution and its relationship to cognition. Xanomeline came out of a development programme aimed at Alzheimer's disease, where the antipsychotic signal was noticed before it was pursued in schizophrenia; the cholinergic tolerability problem that limited it then is what the second molecule was added to solve.

Be exact about what is licensed and what is merely explanatory. The label text held here calls xanomeline a muscarinic agonist and names no receptor subtype at all. The M1 and M4 account is pharmacology and belongs to the rationale, not the licence. Teach it as the rationale; writing it as though the regulator had endorsed a subtype-selective claim upgrades a hypothesis into a licence.

The strategy sits beside the dopamine-blockade algorithm rather than inside it, because nothing in the record read here places it at a line of therapy, ranks it against any comparator, or states after how many failures it is to be tried: no comparative efficacy claim, no sequencing claim and no head-to-head result is held for this agent in this build, and inventing one is more costly than admitting the absence.

GOOD TO REMEMBER

Acute dystonia, parkinsonism and hyperprolactinaemia are consequences of dopamine blockade, so a treatment that does not block dopamine has no mechanistic route to produce them. That is a prediction from mechanism rather than a demonstrated comparative safety finding, and what cannot be said without a trial is any figure comparing rates. Write the mechanism and stop; do not smuggle in a number.

6.3 The peripheral antagonist: what it buys and what it costs

An unopposed muscarinic agonist in the periphery produces a recognisable picture, the reason this class was hard to develop for decades. Predictable from the receptor alone:

- Gastrointestinal stimulation: nausea, vomiting, dyspepsia and cramping, with looser stools.
- Exocrine stimulation: excessive salivation, sweating and lacrimation, and increased bronchial secretion.
- Cardiovascular slowing through vagal muscarinic tone: bradycardia and the symptoms that follow a fall in rate.

Trospium is in the capsule to blunt exactly that list. The cost is that it brings its own peripheral antimuscarinic burden: dry mouth, constipation, urinary hesitancy or retention, blurred near vision, tachycardia, and reduced sweating with its implications in Indian summer heat. The net peripheral picture is therefore a balance of two opposing drugs given together, and the same product can present with either signature: nausea and salivation where the agonist is winning in that tissue, constipation and a dry mouth where the antagonist is. A candidate who has memorised one list will misattribute half the presentations.

The central compartment is deliberately left one-sided, and that asymmetry must be protected. A centrally acting anticholinergic added alongside, whether prescribed for extrapyramidal symptoms out of habit or carried in as a sedating antihistamine, an antispasmodic or a tricyclic, opposes the agonism in the compartment where the treatment must act, and reintroduces the central anticholinergic cognitive load a muscarinic strategy exists to avoid. The anticholinergic burden calculation is not optional here; it is part of the mechanism.

PITFALL

Describing trospium as though it were benztropine or trihexyphenidyl. Those are tertiary amines chosen because they enter the brain, to oppose central cholinergic tone. Trospium is chosen for the opposite property. Calling it an anticholinergic without saying peripherally restricted loses the point of the combination, and is the single sentence most likely to lose the mechanism mark here. The related error is prescribing a central anticholinergic alongside the product.

6.4 The United States: the one grounded record, read field by field

Of the three jurisdictions, exactly one carries a grounded status row, a model of a status claim written properly. The indication, verbatim: COBENFY is indicated for the treatment of schizophrenia in adults. The source is named and narrow: the status was read from the openFDA

Drugs@FDA interface, application NDA216158, with the current structured product label effective 16 January 2026, and the reading was true as of 1 August 2026. Jurisdiction, indication in the source's words, named source, date: the sentence becomes checkable by someone who does not trust you.

Then the field easiest to misreport. The earliest approval-status date in that application record is **26 September 2024**, and the field is named for what it holds, not for what a reader would like it to mean. Whether that date is a first marketing approval, a supplement or something else is a question the record read here does not answer, and this ledger has already published one drug's generic approval date under a first-approval field name and had to retract it. Label a date with the name of the field it came from, every time.

The record carries a caution about itself, which a careful answer respects: openFDA is an agency-operated mirror of Drugs@FDA and of the label archive rather than the Drugs@FDA application itself, and the indication wording is current label text as filed, which is not a statement about any other jurisdiction and must never be carried across one.

This row was tested rather than trusted. The query returned exactly one application, NDA216158, so the row does not blend several applications into one agent-level claim; and three text instruments were run against the primary document, checking that the row text is verbatim, that it is not the primary text cut short and closed with a full stop so that nothing signals the loss, and that the count of indications carried matches the count in the document. All three passed: 262 characters on the row against 262 in the primary, and one indication in the document against one on the row.

1 of 1	0 of 2730	0 of 66	0 of 121
INDICATIONS CARRIED, US ROW VERSUS LABEL	EU REGISTER MEDICINE ROWS MATCHING	CDSCO APPROVAL LISTS WITH AN ENTRY	EXTRACTED TEXTS NAMING TROSPIUM

TRAP **Widening the indication because the mechanism sounds as though it should be wider.** The licensed wording read here is **the treatment of schizophrenia in adults** and nothing else. It is not an indication for negative symptoms as a separately licensed target, for cognitive impairment, for psychosis in dementia, for adolescents, or for use as an add-on to another antipsychotic. The M1 cognitive rationale in 6.2 is what tempts a candidate to widen the sentence, and the widening is unlicensed. Quote the indication whole and let the mechanism stay a mechanism.

6.5 The European Union: an absence in a register that records refusals too

The European position is a declared gap, and a gap's strength depends entirely on the instrument that produced it. The instrument: the European Medicines Agency's own medicines register, generated by the agency on 1 August 2026, downloaded whole and scanned across all 39 columns of all 2730 medicine rows; this substance returns zero rows under any of its names.

What makes that more than a shrug is the range of outcomes the register carries. It records Authorised, Refused, Application withdrawn, Withdrawn, Lapsed, Expired, Revoked, Suspended and Opinion outcomes, so the absence means no central application of any kind was ever

concluded, including one refused or abandoned. The missing document is an agency public assessment report for the substance, or a European Commission implementing decision. Nothing else closes it, and the phrase approved internationally is not a permitted substitute for a named regulator.

Now the limit, where an honest answer separates itself. This register covers the European Union **centralised procedure** only. A product authorised nationally in one or more member states through the decentralised or mutual-recognition procedure would never appear in it. That limit is stated on every European gap row in this build, but it does not bite equally on all of them, and on this agent it is graded as biting strongly: this is a small molecule outside the mandatory centralised scope, so a national authorisation somewhere in the Union is genuinely possible and is wholly unmeasured here. Two further weaknesses are recorded and travel with the finding: the confirming second instrument was run on the grounded rows and on none of the gap rows, so no European gap in this build has independent corroboration, and the assessment-report page would in any case be a second rendering of the same source rather than an independent one, since the register is generated from those pages.

The defensible European sentence is therefore narrow, and worth writing in full: no European Union centralised procedure of any kind is on record for this active substance as of 1 August 2026, which does not exclude a national authorisation that this instrument cannot see.

6.6 India: the affirmative documents, and the sentence that was corrected away

The Indian entry is the one an Indian examination cares about, and it is the one where a confident wrong sentence is easiest to write. The row states that no CDSCO marketing approval was located and that the file is documented at pre-approval stage. That second clause is worth learning, because unlike most gaps in this chapter it is built on documents that exist rather than on documents that do not.

What was found is specific and dated. Two CDSCO **clinical-trial no-objection certificates** name Trospium Chloride plus Xanomeline, one against an IQVIA file numbered 44439 and one against a Bristol-Myers Squibb file numbered 52364. Minutes of the **subject expert committee** for neurology and psychiatry dated 30 May 2025, 10 December 2025 and 22 December 2025 record Bristol-Myers Squibb India Private Limited presenting the Phase III protocols CN0120051 and CN0120025, with the committee recommending permission to conduct the trial.

That footprint licenses one sentence and forbids another. A trial permission is not a **marketing approval**. The authority that would settle the question is the Central Drugs Standard Control Organisation and the Drugs Controller General of India, and expert committee minutes, clinical-trial no-objection certificates, International Cell written confirmations, CTRI registration, brand listings, pharmacy availability and manufacturer copy are none of them marketing authorisation. A **CTRI registration** deserves separate mention because it is the substitution an Indian candidate is likeliest to make: it records that a trial exists, not permission to market. What would

close the row is an affirmative CDSCO or DCGI new-drug permission, a current CDSCO-approved label, or an entry in an official CDSCO approval list.

Note what the row does not say, because it was deliberately taken out. The lead sentence of this row once read that the agent was not approved in India at retrieval, and that was corrected: pre-approval activity by two named firms is evidence that those firms had not yet been granted permission, and it is not proof that no other holder exists, in a corpus this build itself demonstrates to be incomplete. The evidence did not change; only the claim it was being made to carry was narrowed, and the row remains a declared gap either way. No India row in this build asserts that a drug is not approved in India. Each asserts that no CDSCO approval record was located, and names what was searched.

IN INDIA

Write it in this order. No CDSCO marketing approval record for xanomeline with trospium was located as of 1 August 2026; the Indian file is documented at clinical-trial stage, with two trial no-objection certificates and subject expert committee minutes of 2025 recording Phase III protocols and a recommendation to permit the trial; a trial permission is not a marketing approval. Do not write that the drug is not approved in India, which is stronger than this evidence supports. Do not write that it is approved, because nothing retrieved says so. For a patient or a family: the treatment is not something an Indian service can offer today outside a trial, and a sponsor's trial permission is not an availability promise to a patient.

6.7 The instrument that cannot see half of this drug

The India row carries one clause that looks like evidence and is worth nothing, and why is the most transferable thing on this item: the agent is absent from all 66 published CDSCO approval lists.

Trospium is absent from all 66 of those documents and from all 121 extracted text files under every spelling tried, including variants tolerant of the rn and m ligature confusion, of 0 against o and 1 against l, and of letter-spaced layout, and under the brand and programme names. It is one of three drugs that failed as **positive controls** for this corpus, alongside agomelatine and molnupiravir: drugs taken to be approved and marketed in India that the corpus nonetheless cannot display. The corpus has been shown incapable of displaying trospium at all, so its silence is a property of the instrument and not a fact about the drug.

The row records this explicitly: the absent-from-66-lists clause is worth nothing for trospium specifically and must not be read as a negative finding about it, and the verdict does not depend on that clause, because the verdict rests on the affirmative documents, the two trial no-objection certificates and the dated committee minutes, which are unaffected. The template: a null reading from an instrument shown blind to the target contributes nothing, and the same page can still carry a sound conclusion resting on something else.

Two sensitivity facts belong to any Indian gap sentence on a recent agent. Every drug this corpus is known to find sits in a document dated 2023 or earlier, and independent sensitivity in the 2024 to 2026 window is 0 of 1, so the corpus has never been shown to find a recent approval;

and the positive-control tally of 17 of 20 may not be cited as warrant for a gap on any agent approved after 2023. The corollary is that a gap is a statement about a search and not about a product, which is also why the three text-comparison instruments were recorded as not applicable to both gap rows on this agent: a declared gap carries no indication text, so an instrument that compares text against a primary document has nothing to test.

TRAP **Quoting a null result from an instrument that was never shown able to detect the thing.** Placed beside two genuine documents, the sentence that trospium appears in none of 66 approval lists reads as a third piece of evidence, and is not evidence at all. The test before quoting any absence: was this instrument ever demonstrated to return a known positive? Here the answer is no, recorded as one of three control failures. The mirror-image error is to conclude the whole row collapses because one clause is empty; the verdict was built on the affirmative documents and never on the absence.

6.8 The three records side by side, and what a complete answer contains

The same drug, the same date, three records of entirely different kinds. Reproduce the differences and the question is answered in any form.

JURISDICTION	WHAT THE RECORD HOLDS	THE SENTENCE THAT MAY BE WRITTEN	THE SENTENCE THAT MAY NOT
United States	A current label indication for the treatment of schizophrenia in adults, application NDA216158, read 1 August 2026	The indication quoted whole, with the application, label date and reading date	Any wider indication, and any date labelled as a first approval
European Union	Zero rows in the agency medicines register of 2730 medicines, which also records refusals and withdrawals	No centralised procedure of any kind is on record, as of the reading date	Not authorised in Europe, which the centralised-only instrument cannot support
India	Trial no-objection certificates and 2025 committee minutes recording Phase III protocols and a recommendation to permit, and no located approval	No CDSCO approval record located; the file is documented at trial stage	Not approved in India, and any negative inference from the approval-list corpus about trospium

Reduced to four sentences, the item is this.

- The mechanism is muscarinic, not dopaminergic; the combination works because a permanently charged antagonist stays peripheral while the agonist acts centrally.
- The adverse-effect picture balances two opposing peripheral actions, so either a cholinergic or an antimuscarinic presentation is possible from one product, and a centrally acting anticholinergic added alongside attacks the mechanism.
- The regulatory position has three answers and no single one, each carrying its jurisdiction, its indication in the source's words, its named source and the date it was read.

- The Indian file is a declared gap standing on affirmative documents, and its one clause reporting an absence is void because the instrument is blind to trospium.

MNEMONIC

Charge, Compartment, Country, Control

Charge is why trospium stays out of the brain. Compartment is where each molecule is meant to act, and why the pairing is not self-cancelling. Country is the jurisdiction every status sentence must name and never leave. Control: before quoting any absence, ask whether the instrument was ever shown able to find what it was looking for.

The antagonist is in the capsule to be kept out of the brain, the approval is in one country and must stay there, and an absence reported by a blind instrument is not a finding.

7 · Endoxifen and protein-kinase-C-targeted agents in mania

This is the one agent in the chapter with an affirmative Indian record and a measured American absence, and the examinable content of the item is that both sentences are true at once. The marks sit on the regulatory question: writing about one molecule in three jurisdictions without letting any of them leak into another. It goes four ways. Writing "approved" with no country attached. Converting an American search that found nothing into a statement that nothing exists. Reading a tablet strength printed inside a product-name cell as a dose. And quoting the Indian indication with its qualifier quietly missing, the defect this build planted on this entry to prove its instruments could catch it. The acute-mania treatment algorithm is not rebuilt here.

7.1 A metabolite promoted to a drug, and what that changes

Endoxifen is an **active metabolite** of tamoxifen, and tamoxifen is named here only as the parent compound. When a parent owes its activity to a metabolite, response depends on a conversion step, and that step is where people differ: genotype, competing inhibitors and adherence all act on it before any target is reached. Give the metabolite itself and the step leaves the chain, taking its variability with it.

That is an argument about pharmacological architecture and not about magnitude. It says where variability enters, not how much of it there is, and no exposure figure, potency ratio, conversion proportion or genotype frequency is held anywhere in the material behind this chapter. It is also a different drug from its parent for regulatory purposes, with separate applications, indications and approval records, so nothing true of tamoxifen may be carried onto endoxifen in either direction.

GOOD TO REMEMBER

Giving a metabolite removes a source of variation between patients. It does not remove the need to show that the metabolite works. Removing a conversion step is a pharmacokinetic argument and the prescriber's question is a pharmacodynamic one, so answering the second with the first answers the wrong question.

7.2 The protein kinase C rationale, and what kind of claim a rationale is

Protein kinase C is not a receptor. It is a family of intracellular signalling enzymes downstream of receptor activation, acting on the machinery that changes what a cell does over hours and days rather than milliseconds. The intervention is aimed at a node inside the cell, several steps from the synapse, and the phenotype it must move is a whole manic episode.

The shape of that proposal is what an examiner can test. It is target-first: start from a mechanistic node, find a compound with activity there, and predict a clinical effect in a syndrome the node is thought to join. Most older agents arrived the other way round, from a clinical observation worked backwards. So its strength and its weakness are one feature: it is derived rather than observed.

A rationale licenses a hypothesis and a trial. It does not license a statement about efficacy, about how quickly an effect appears, about how large it is, or about how the agent stands against any existing treatment. No artefact behind this chapter holds trial evidence for this agent: no trial name, no sample size, no endpoint, no effect size, no response rate, no dose, no titration step, no adverse-event rate. This account names none of them, and offers no comparison against any existing mood stabiliser, because none is held.

RULE **A mechanistic rationale is a reason to run a trial and never a substitute for one.** Could your sentence be false even if the mechanism is right? "Protein kinase C is implicated in mania" and "inhibiting protein kinase C treats mania" are separated by the whole of clinical trial evidence.

7.3 India: the affirmative record, read field by field

The Indian record is the strongest of the three here, which reverses the order candidates expect. Its source is the published new-drug approval lists of the Central Drugs Standard Control Organisation, CDSCO: 66 distinct documents deduplicated by content hash, downloaded direct over verified transport on 1 August 2026, every one yielding an extractable text layer so that no absence in the corpus is an unreadable-file artefact; the entry is serial number 20 in the list file ndedec13appro.pdf.

The indication, quoted whole because a paraphrase is where the qualifier goes missing: For the acute treatment of manic episodes with or without mixed features of Bipolar I disorder. Endoxifen citrate bulk and Endoxifen tablets 8 mg. The first leg is an indication bounded three ways: acute treatment, manic episodes, bipolar I. The words **mixed features** are the part most often dropped, and dropping them changes who the indication covers. The second is a product-name leg naming a **bulk drug** entry and a tablet strength. That strength is a formulation fact

printed inside a name cell: not a dose, not a starting dose, not a maximum, and no dosing information for this agent exists in any artefact behind this chapter.

Then the date, the field this build has already had to correct once. The entry carries **10 December 2019**, and the row now describes that as the quantity printed in the column headed **Date of issue** rather than as an approval date, the earlier wording having claimed to quote a label the document does not print. The reading is dated too: the status true as of date on this row is 1 August 2026.

Corroboration is thinner than the verbatim quotation looks. It rests on one CDSCO document. Date and indication were confirmed against the raw column layout of the source table rather than its flattened text, and the neighbouring date of 20 November 2019 was checked and belongs to serial number 19 and to an entirely different substance. In a printed table a date can be read off the row above with nothing looking wrong.

An older record of the same search read this jurisdiction as a gap and named what would close it: a CDSCO or DCGI permission, an approved label, or an official approval-list entry with its date and indication verbatim; and what would not: a CTRI registration, a brand listing, pharmacy availability or manufacturer copy. What closed it was the third permitted item and nothing weaker. The question is never whether a document exists, but whether it can settle the question.

IN INDIA

Write it in this order. An entry for endoxifen appears in a CDSCO published new-drug approval list at serial number 20, carrying the indication above and the date 10 December 2019 under the printed heading Date of issue, true as of 1 August 2026. Quote the indication with its mixed-features qualifier intact. Do not turn the 8 mg tablet strength into a dosing instruction. And be clear what the record does not settle: it is an approval-list entry, not a statement about supply, price or whether your pharmacy stocks it.

7.4 How that Indian cell was certified, and the defect planted in it

Quoted correctly and read correctly are two achievements; this row was tested for the second. Three text instruments were run against the primary document: one checking that the row text is verbatim and that the source sentence names the primary, one checking that the row is not the primary cut short, at 146 characters against 146, and one checking that the indication cell is carried in full, the primary cell measuring 94 characters. All three passed against a locator naming the document, page 4, serial number 20, the bounds of the indication column and the row band, with corroboration recorded as trustworthy.

The rule behind that verdict: a cell counts as read when a reader working from the drawn rules of the table agrees with at least one independent reader, and two readers of the same flattened-text kind are not independent, because they share whatever made the first of them wrong.

Now the part that makes this row the most instructive object in the section. This entry is the one the build used to test its own instruments, planting two defects in it. The first cut the indication at a clause boundary and closed it with a full stop where the primary entry continues with or

without mixed features of Bipolar I disorder, dropping a qualifier that changes who the indication covers. The second removed the same qualifier from the middle of the cell and left the ending intact, so nothing was truncated at the end: that plant exists to separate the two instruments, because the prefix instrument must stay silent on it and only the completeness instrument may fire. Both guards fired and both plants were caught.

Note what both plants attacked. It was the clinical qualifier.

TRAP **Quoting an indication that ends in a full stop and assuming it ended there in the source.** Cut at its clause boundary, the Indian entry reads as a complete, plausible sentence with the mixed-features qualifier gone, and a cut closing with valid punctuation leaves no mark, so only comparison against the primary finds it. Reproducing an approved indication means reproducing its qualifiers: population, phase, severity, specifier. And a gate never shown to catch the defect it was built for has not been tested; planting the defect is the test.

3 of 3	0 of 5	228	0 of 2730
INDIA TEXT INSTRUMENTS PASSING	US QUERY ROUTES RETURNING A RECORD	LABEL DOCUMENTS MATCHING THE STRING	EU REGISTER MEDICINE ROWS MATCHING

7.5 The United States: a measured absence, and the sentence it does not license

An absence that was measured is a different object from an absence that was assumed, and only one of the two may be written down. A query of the Drugs@FDA interface returned no application record for the generic name endoxifen on 1 August 2026, the endpoint answering with its own not-found response. That probe found nothing; it does not establish that the agent is unapproved in the United States. To close the row a human must read Drugs@FDA directly and, if there is no application, say so on the agency's authority rather than on this probe's silence. The authority any American claim would rest on is the United States Food and Drug Administration, through Drugs@FDA.

The gap was then strengthened into a **measured absence**. Five query routes were run, every one returned the endpoint's own not-found response, there were zero transport failures, and every route was positive-controlled in the same run against a value known to return records, so a zero here is a route that answered and found nothing, not a route that failed. That **positive control** is the difference: without one, a broken route and an empty answer look identical from outside. The five routes were four fields of the application interface, generic name, substance name, active ingredient name and brand name, and one field of the label archive.

The scope is narrower than it first reads. It measures the Drugs@FDA interface and the label archive only. It is not the agency's own affirmative statement that no application exists, and it is silent on investigational and breakthrough designations, which appear in neither instrument. So the defensible American sentence names the search and stops. Writing it harder is forbidden, because an American gap in this build has already been wrong once: one agent's record carries no cross-reference block at all, so three of five routes returned not-found on it while the active-ingredient route returned the record.

Then the figure showing how a search can look conclusive in the wrong direction. A free-text query of the label archive for the string endoxifen, run during independent adjudication on 2 August 2026, returned a successful response with 228 matching documents, while every structured route above returned zero. The two are not in conflict. A **free-text search** matches a string anywhere in a document, so a document belonging to another product that mentions the substance counts as a hit, whereas a structured field query asks whether the substance is the product's own named ingredient.

TRAP **Collapsing an Indian approval-list entry and an American gap into one sentence about the drug.** This agent is the border-crossing error waiting to happen, both ways. Writing that endoxifen is approved, with no country attached, exports an Indian entry to the United States and Europe. Writing that it is not approved in the United States converts a search that found nothing into a regulator's negative finding, which no instrument here supports. Three jurisdictions, three records, three sentences, each with its own source and date.

7.6 The European Union: a register that records refusals, and the limit that follows

The European position is also a declared gap, and a gap is only as strong as the instrument behind it. The instrument is the European Medicines Agency's own published medicines register, generated by the agency on 1 August 2026, downloaded whole and scanned across all 39 columns of all 2730 medicine rows; this substance returns zero rows under any of its names.

What lifts that above a shrug is the range of outcomes the register carries. It records authorised, refused, application withdrawn, withdrawn, lapsed, expired, revoked, suspended and opinion outcomes, so an absence means no central application of any kind was ever concluded, including one refused or abandoned. The missing document is an agency public assessment report for the substance, or a European Commission implementing decision, and the phrase approved internationally is not a permitted substitute for a named regulator.

Now the limit. The register covers the European Union **centralised procedure** only, so a product authorised nationally through the decentralised or mutual-recognition procedure would never appear in it. That limit is printed identically on every European gap row, so it is graded per agent. For this substance it is graded as biting moderately: a small molecule outside the mandatory centralised scope could genuinely hold a national authorisation this instrument cannot see. The confirming instrument ran on all eight grounded European rows and none of the eight gap rows, so no European gap here has independent corroboration, and the assessment-report page would anyway be a second rendering of the same source. Second-instrument coverage of European gaps is **0 of 8**.

So the defensible European sentence is narrow: no European Union centralised procedure of any kind is on record for this substance as of 1 August 2026, which does not exclude a national authorisation this instrument cannot see. The three text instruments used on the Indian entry report nothing here, because a declared gap carries no indication text: a gap is not a claim about a product, it is a claim about a search. Each record was built in isolation, the Indian row consulting no American or European value, and the European row no American or Indian one.

7.7 The three records side by side, and what a complete answer contains

One molecule, one reading date, three records of different kinds.

JURISDICTION	WHAT THE RECORD HOLDS	THE SENTENCE THAT MAY BE WRITTEN	THE SENTENCE THAT MAY NOT
India	An approval-list entry at serial number 20 carrying the acute manic episode indication with or without mixed features, a bulk-drug leg and a tablet strength, dated 10 December 2019 under the heading Date of issue	The indication whole, with the document, serial number, date heading and reading date	Any dose built from the strength in the name cell, or any claim about supply
United States	Five positive-controlled query routes returning the endpoint's own not-found response, zero transport failures	These named routes, on this date, returned no application record	Not approved in the United States, which converts a probe's silence into the regulator's finding
European Union	Zero rows in an agency register of 2730 medicines that also records refusals and revocations	No centralised procedure of any kind is on record at the reading date	Not authorised in Europe, which a centralised-only instrument cannot support

Reduced to five sentences, the item is this.

- Three jurisdictions, three records, three sentences: an approval-list entry, a measured absence and a register absence, each with its own source and date.
- The Indian indication is quoted whole because the mixed-features qualifier decides who it covers, and the strength beside it is a formulation fact, not a dose.
- A protein kinase C target is a rationale for a trial, not a result from one, and no trial figure is held.
- An absence is a statement about a search: name the routes, their controls, and what the instrument cannot see.
- A date is only as strong as its field name, and a printed column heading is not a regulatory event.

MNEMONIC

Metabolite, Mechanism, Mixed features, Measured

Metabolite is what the drug is. Mechanism is a rationale until a trial says otherwise. Mixed features is the qualifier that must survive into your answer. Measured is the only kind of absence you may write down.

One molecule, three records: an Indian approval-list entry quoted whole with its qualifier, an American absence that was measured rather than assumed, and a European silence a national authorisation could still be hiding behind.

8 · Anti-amyloid monoclonal antibodies as a recent advance (cross-reference to dementia)

This item is a pointer, and knowing why it is a pointer is itself examinable. **Anti-amyloid monoclonal antibodies** reach a recent-advances list because they changed what a therapeutic claim here may sound like, and they sit in another chapter because their mechanism, trial evidence, imaging complication, genotype caveat, selection criteria, monitoring and regulatory record are taught with the Dementias and neurocognitive disorders notes. What is left is what this chapter owns: the class of claim these agents exemplify, the shape such an argument must take, why the regulatory position is not printed here, and what an Indian service must decide.

8.1 Why this is a pointer, and what a pointer owes the reader

This material is taught with the dementia notes, and the item itself is worded as a recent advance carrying a cross-reference there, so a pointer is what is being asked for and the reason is not space. Two accounts of one drug class in one paper are written at different times against different sources and drift apart, and the drift stays invisible until a candidate quotes the thinner one and is corrected out of the fuller one.

A pointer owes the destination named exactly, some idea of what is there, and a plain statement of what is withheld. The owning chapter carries dedicated passages on each of these agents and on the imaging complication, so the material is there and at length. Everything listed at the head of this section is taught with the Dementias and neurocognitive disorders notes and none of it is restated here. Going quiet on all of that without saying so invites a reader to fill the silence from memory.

8.2 The class of claim, and why psychiatry has an interest in it

Almost everything a psychiatrist prescribes makes a **symptomatic** claim: the current state will be better while the drug is taken, and it is judged on that state. A **disease-modifying** claim differs in kind: it asserts an effect on the course of the illness rather than on the state at a point along it, so it cannot be demonstrated at a point. It needs an observation period long enough for two trajectories to separate, an agreed way of locating a patient on one, and an argument that the gap between them matters to the person living it. Each is a place the claim can fail without the drug doing anything wrong.

Which is why the class belongs in a psychiatry chapter: these patients reach **old-age psychiatry** clinics as well as neurology ones, and the first clinician a family asks is often a psychiatrist. And the argument travels, because the way a disease-modifying claim is first pressed becomes the template the next class is held to.

RULE **A symptomatic claim is judged on the state, a disease-modifying claim on the course, and the same evidence cannot settle both.** Decide which claim is on trial before appraising any result: a better state at a point is silent on the course, and separating trajectories show nobody better today.

8.3 The shape of the evidence argument, with no figure from it

Such an argument runs at two levels and keeping them apart is where the marks are. **Target engagement** is the first: the thing the drug acts on has changed. Patient benefit is the second: the person does better for longer. Engagement alone is a **surrogate marker** argument, only as strong as the assumption joining marker to outcome, and the history of therapeutics is largely a history of that assumption failing. A marker result is a reason to keep testing, not to treat.

At the patient level the problem is what a difference on a rating scale means. A difference can be reliably present and too small for anyone to notice, and the bridge is the **minimally important difference**, the change below which no patient or family could tell. Where that threshold sits is contested rather than settled, so one honest appraiser reads a result as a worthwhile gain and another as clinically negligible; state both readings and name the threshold each rests on. Two constraints join the balance: a group difference in an average rate of change is not a prognosis for the individual in front of you, and where a treatment carries a delivery burden and a surveillance requirement the benefit side must weigh more before the balance tips.

No figure from that argument appears here, by rule and not oversight. The chapter holds no grounded evidence lane for this item: no trial name, no sample size, no observation period, no effect size, no adverse-event rate. Writing one from memory is a fabricated citation, the one defect that cannot be repaired later.

8.4 Why the regulatory position is not on this page, and what a delegated gap is

The regulatory position of these agents is not written here, and the reason earns more marks than the position would. The chapter's status ledger holds **48** rows and not one is an anti-amyloid row. That is neither an oversight nor a failed search: the ledger says no row is written there deliberately, so that none can be quoted from there, and it names the three agents covered, lecanemab, donanemab and aducanumab.

Three silences look identical here and mean different things. A searched gap is a question a named instrument was pointed at, which returned nothing; the honest sentence names what was searched and whose authority a claim would rest on, and never hardens into a categorical negative about a jurisdiction. An instrument blindness is a question the instrument could not have answered whatever the truth was, so its silence carries no information. A **delegated gap** is neither: the fact exists, another owner holds it, and this page declines to restate it so two accounts of one fact cannot diverge. This item is the third kind, and the only one that is not an absence of fact.

A delegation is not a warrant, though: nothing read for this chapter establishes the evidentiary state of the owner's account, which predates this chapter's four-field standard and carries its own warning against giving its record false precision. So go to the owner for what it holds, read

anything found there against the jurisdiction it was written about, and never carry a position across a border. Until that chapter is refreshed this class has a delegated status, not an unknown one.

0 of 48	3	131	16
CHAPTER STATUS ROWS ON THESE AGENTS	AGENTS WHOSE STATUS ANOTHER CHAPTER HOLDS	OWNER PASSAGES MATCHING LECANEMAB, DONANEMAB OR THE IMAGING COMPLICATION	OTHER AGENTS THIS CHAPTER HOLDS A STATUS FOR

TRAP **Supplying the regulatory position from memory because the section did not.** A reader given no jurisdiction, no wording and no date often finishes the sentence privately, and a half-remembered headline surfaces attached to the wrong country and year. A status sentence is worth writing only when it carries the jurisdiction, the wording the regulator used, the source it was read from and the date it was true; none of the four is available here for these agents.

8.5 What an Indian service is actually being asked to decide

The decision facing an Indian psychiatrist is not whether to prescribe one of these agents, and treating it as though it were is how the topic gets answered badly. Four decisions are live now, whatever any regulator has done. What to say to a family who arrive having read something: name what kind of treatment is discussed, say what is held elsewhere and can be checked, and offer to find out rather than guess, because a warm guess is repeated as though a doctor confirmed it. Whether raising the class unprompted is kind where any pathway would be long and uncertain, which is a judgment about this family and not a rule. What would have to exist before the question became a real option, which is the pathway set out below. And affordability, which in Indian practice usually settles the question first.

IN INDIA

Write the Indian answer in this order. The regulatory position is not stated here and is not to be inferred from the silence. Nothing read for this chapter holds an Indian position for these agents: the status ledger writes no row for them at all, and the material with the Dementias and neurocognitive disorders notes is written about the United States and Europe, its own gap list recording that current Indian regulatory, availability and cost information is not supplied there. So nothing found at the owner may be read across into an Indian answer. What an Indian old-age service must establish for itself is the pathway: who confirms the diagnosis, where the investigations are done, who owns the decision, who carries the follow-up, and what the family can afford.

8.6 The boundary in both directions, and what a complete answer contains

QUESTION	WHERE IT IS HELD	WHAT THIS PAGE MAY SAY
How the antibodies act, and on what	The owning item	Only that the class name states its target
What the trials showed, in numbers	The owning item	Nothing numerical: no evidence lane exists here
Imaging complication, genotype caveat, selection, monitoring	The owning item	Only that they exist and dominate the safety discussion
Regulatory position in any jurisdiction	The owning item, a delegated gap until refreshed	Nothing, for any agent, in any jurisdiction, on any date
The class of claim, and what a service must decide	Here	Claim type, argument shape, the delegation, the Indian pathway

Reduced to four sentences:

- A pointer by design: two accounts of one drug class diverge, and the thinner loses marks.
- The claim is disease-modifying, not symptomatic, which changes what evidence could settle it and how long it must run.
- The argument has two levels; a marker result licenses no patient conclusion, and the noticeability threshold is contested, so both readings are stated.
- The status is a delegated gap, not an unknown: held by its owner, restated nowhere here, and not yet meeting the four fields a status sentence needs.

GOOD TO REMEMBER

The commonest real-world error here is conversational, not pharmacological: a family told kindly and vaguely that something new exists and might help leaves having heard a promise nobody made. Name what is known, what is held elsewhere, and what you will find out.

MNEMONIC

Claim, Course, Custody, Clinic

Claim: symptomatic or disease-modifying, decided first. Course: judged over a trajectory, never at a point. Custody: another owner holds the record; this page restates none of it. Clinic: pathway, follow-up, affordability.

This item is a pointer: the class of claim and the service decision are taught here; mechanism, evidence and regulatory position are quoted from their owner or not at all.

9 · Moncrieff serotonin umbrella review and the serotonin-hypothesis debate

The examinable object here is a piece of evidence appraisal, not a drug and not a neurotransmitter. In 2022 a group of authors led by Joanna Moncrieff published an umbrella

review in Molecular Psychiatry asking whether the research literature supports the idea that depression is caused by low serotonin, reported that it does not, and set off an argument that reached patients before it reached examiners. Every mark in this topic is won or lost on discipline of attribution: what the review found, what its authors concluded, what its critics answered. The serotonin system itself, its pathways and its receptor subtypes, belongs to the Neurochemistry, Neuroendocrinology, Neuroimaging and Basic Psychopharmacology notes, and the monoamine hypothesis of depression and its limitations belong to the Mood disorders: classification, depressive and bipolar syndromes, neurobiology notes.

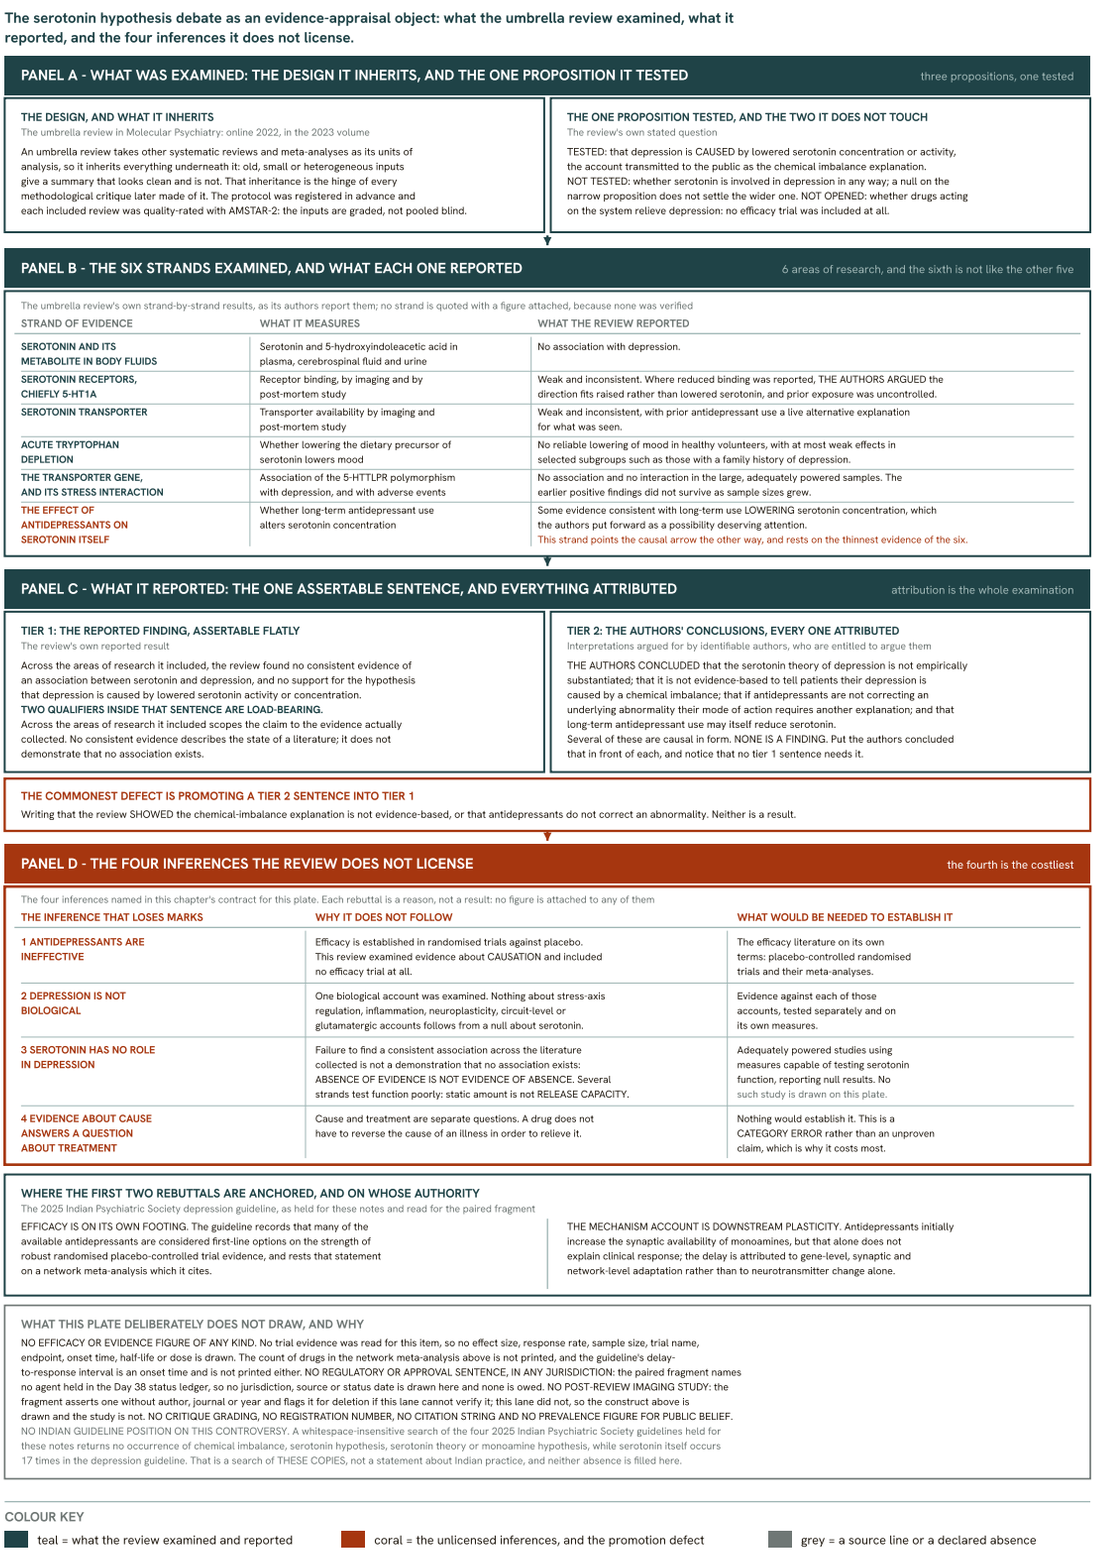


Fig 38.5 The serotonin hypothesis debate: what the umbrella review examined, what it reported, and the four inferences it does not license. One narrow proposition was tested, that depression is caused by lowered serotonin concentration or activity; neither the wider serotonin question nor efficacy was opened, because no efficacy trial was included. Six strands carry a reported finding and no figure, the sixth pointing the causal arrow the other way. What it reported divides into the finding, assertable flatly with both its qualifiers and the authors' conclusions, every one attributed. (The 2022 umbrella review in Molecular Psychiatry; the 2025 Indian Psychiatric Society depression guideline as held for these notes.)

9.1 What an umbrella review is, and why the design carries the whole argument

An **umbrella review** is a systematic review whose units of analysis are other systematic reviews and meta-analyses rather than primary studies. It sits at the top of the hierarchy of evidence and inherits everything underneath it. If the reviews it collects are old, small, heterogeneous or poorly conducted, so is the summary, however clean it looks. That inheritance is the hinge of every methodological criticism later made of this paper.

One housekeeping point before the method. The paper appeared online in 2022 and was carried in the journal's 2023 volume, so both years are correct and a marker who sees the other one has not caught an error.

Three features of the method matter for an answer. The protocol was registered in advance, which matters because a review of a contested hypothesis is exactly the kind whose scope can drift toward the answer its authors expect. The quality of each included review was rated with **AMSTAR-2**, a standard appraisal instrument for systematic reviews, so the paper does grade its own inputs rather than pooling them blind. And where no systematic review existed for a strand, the authors went to large individual studies instead. That decision is defensible, and also a departure from the pure design critics later used against the paper.

The review did not ask whether serotonin is involved in depression in any way. It asked whether the evidence supports a much narrower proposition: that depression is caused by lowered serotonin concentration or activity, the account transmitted to the public as the **chemical imbalance** explanation. Those are different propositions, and almost every error in this topic is a slide between them.

RULE **One narrow proposition was tested, and every sentence written about the review must name which proposition it is about.** "Depression is caused by low serotonin" is not the same claim as "serotonin is involved in depression", which is not the same claim as "drugs acting on the serotonin system relieve depression". The first was tested. The second was not tested and is not settled by a null result on the first. The third belongs to a different literature that this review did not open.

9.2 The six strands, and what each one reported

The evidence was organised into **six** areas of serotonin research, and the organisation is itself examinable: a question asking what the review examined is asking for this list. Reproduce it with a finding attached to each line, never as a bare list.

STRAND OF EVIDENCE	WHAT IT MEASURES	WHAT THE REVIEW REPORTED
Serotonin and its metabolite in body fluids	Concentrations of serotonin and of 5-hydroxyindoleacetic acid in plasma, cerebrospinal fluid and urine, compared between depressed and non-depressed groups	No association with depression
Serotonin receptors, chiefly the 5-HT1A receptor	Receptor binding, by imaging and by post-mortem study	Weak and inconsistent. Where reduced binding was reported, the authors argued the direction fits raised rather than lowered serotonin, and prior antidepressant exposure was not controlled for
Serotonin transporter	Transporter availability by imaging and post-mortem study, and its relation to depression	Weak and inconsistent, with prior antidepressant use a live alternative explanation for what was seen
Acute tryptophan depletion	Whether experimentally lowering the dietary precursor of serotonin lowers mood	No reliable lowering of mood in healthy volunteers, with at most weak effects in selected subgroups such as those with a family history of depression
The serotonin transporter gene, and its interaction with stress	Association of the 5-HTTLPR polymorphism with depression, and the proposed interaction between that polymorphism and adverse life events	No association and no interaction in the large, adequately powered samples; the earlier positive findings did not survive as sample sizes grew
The effect of antidepressants on serotonin itself	Whether long-term antidepressant use alters serotonin concentration	Some evidence consistent with long-term use lowering serotonin concentration, which the authors put forward as a possibility deserving attention

Read the sixth row differently from the first five. The first five ask whether a serotonin abnormality precedes depression; the sixth asks whether treatment produces one, which points the causal arrow the other way, and it is the strand carrying the authors' own hypothesis most visibly. It also rests on the thinnest evidence of the six. An answer that reports it with the same confidence as the genetic strand has flattened the paper into a slogan.

9.3 Three tiers: the finding, the authors' reading, and what neither warrants

Attribution is the whole examination here, so build the answer in three labelled tiers. The first tier is the review's reported finding, and it may be asserted flatly: across the areas of research it included, the review found no consistent evidence of an association between serotonin and depression, and no support for the hypothesis that depression is caused by lowered serotonin activity or concentration. Two qualifiers inside that sentence are load-bearing. "Across the areas

of research it included" scopes the claim to the evidence actually collected. "No consistent evidence" describes the state of a literature; it does not demonstrate that no association exists.

The second tier is what the authors concluded, and every sentence of it carries their name. They held that the serotonin theory of depression is not empirically substantiated; that it is not evidence-based to tell patients their depression is caused by a chemical imbalance; that if antidepressants are not correcting an underlying abnormality then their mode of action requires some other explanation; and that long-term antidepressant use may itself reduce serotonin. These are interpretations, several of them causal in form, argued for by identifiable authors who are entitled to argue for them. They are not findings, and a script that presents them as findings has surrendered the marks this topic exists to award.

The third tier is the set of inferences that neither the findings nor the authors' conclusions warrant. Four recur, and an examiner is watching for them.

TRAP **Promoting a second-tier sentence into the first tier.** The commonest version is writing that the review showed the chemical-imbalance explanation is not evidence-based, or that antidepressants do not correct an abnormality. Neither is a result. The result is that a set of literatures did not consistently support one proposition; the two sentences above are what identifiable authors concluded from it, and they were contested in print within months. The fix costs four words: put "the authors concluded that" in front of every sentence in the second tier, and notice that no sentence in the first tier needs it.

THE INFERENCE THAT LOSES MARKS	WHY IT DOES NOT FOLLOW	WHAT WOULD BE NEEDED TO ESTABLISH IT
Antidepressants are ineffective	Efficacy is established in randomised trials against placebo. This review examined evidence about causation and included no efficacy trial at all	The efficacy literature on its own terms: placebo-controlled randomised trials and their meta-analyses
Depression is not biological	One biological account was examined. Nothing about stress-axis regulation, inflammation, neuroplasticity, circuit-level or glutamatergic accounts follows from a null result about serotonin concentration	Evidence against each of those accounts, tested separately and on its own measures
Serotonin has no role in depression	Failure to find a consistent association across the literature collected is not a demonstration that no association exists, and several strands rested on measures that test serotonin function poorly	Adequately powered studies using measures capable of testing serotonin function, reporting null results
Evidence about causation answers a question about treatment	Cause and treatment are separate questions. A drug does not have to reverse the cause of an illness in order to relieve it	Nothing would establish it. This is a category error rather than an unproven claim, which is why it is the costliest of the four

6	2022	1	4
AREAS OF RESEARCH EXAMINED	YEAR OF PUBLICATION	FINDING ASSERTABLE AS THE REVIEW'S OWN	INFERENCES THAT LOSE MARKS

9.4 Why evidence about cause cannot answer a question about treatment

The mechanism by which an effective drug works is not required to be the reverse of the cause of the illness it treats. Analgesia relieves headache and headache is not caused by a shortage of analgesic. The **corrective model**, in which a drug replaces something the body is missing, describes a small minority of effective treatments rather than the general case, which is why the insulin analogy proves less than it appears to. An aetiological null result and a treatment effect are therefore logically independent, and a paper about the first leaves the second where it found it.

Indian guidance puts efficacy on its own footing, which lets this be answered with a source rather than an argument. The 2025 Indian Psychiatric Society depression guideline records that many of the available antidepressants are considered first-line options on the strength of robust randomised placebo-controlled trial evidence. The guideline rests that statement on a network meta-analysis of **21 antidepressant drugs**, which it cites as having confirmed that all these agents are more effective than placebo. Both rest on trials rather than on any theory of causation, and both would read identically if every strand of the umbrella review had come out the other way.

TRAP **Answering an efficacy question with the umbrella review.** A stem asking whether antidepressants work, or whether this controversy changes prescribing, is testing this boundary, and the review is the wrong instrument. The correct move is to say the question is answered by the efficacy literature, name what that literature consists of, then say in one sentence what the umbrella review does bear on: how the drug is explained to the patient rather than whether it is offered. A candidate who cites a review of aetiological evidence as grounds for a prescribing decision has answered a different question well.

9.5 The critiques, and which of them actually bite

Four objections were published, and are not of equal weight. Grade them rather than listing them.

The first is the straw-man objection: that no well-informed academic psychiatrist held a simple chemical-imbalance theory, so a review dismantling it defeats a position the field had already left. Half of that is true, and the Indian guideline says so in its own voice. Its box on the neurobiological mechanisms of antidepressant action opens by recording that antidepressants initially increase the synaptic availability of monoamines such as serotonin and noradrenaline, but that this alone does not explain clinical response. The other half is where the objection fails: the review's target was not only the academic theory but what patients are told, and that account is widely believed and widely used in consultations. A theory abandoned in the seminar room and still spoken in the clinic is a fair thing to examine.

The second is the tool objection, and technically the strongest. An umbrella review is a good instrument for a question a coherent body of systematic reviews was built to answer, and a poor one for assembling several distinct literatures, each measuring a different construct by different methods, into a single verdict on one hypothesis. Critics argued that the included reviews were

heterogeneous in age, quality and purpose, that the quality gradings were contestable, and that some strands, the depletion literature in particular, were read less favourably than the primary studies allow. A multi-author reply published in the same journal, led by Jauhar with several academic psychiatrists, argued along these lines that the serotonin system is implicated in depression and that this method was not capable of showing otherwise. Note what that reply does not do: it contests the inference and the instrument, and it does not restore a simple low-serotonin account.

The third is the logical one, and the cleanest to write. **Absence of evidence is not evidence of absence.** A literature that is underpowered, cross-sectional, heterogeneous in its measures and collected over decades can fail to detect a real effect. Converting the review's failure to find consistent association into a demonstration that no association exists requires a step the design cannot supply. This objection is sound, and it cuts at the authors' broader conclusions rather than at their reported finding, which is why the three tiers above matter.

The fourth concerns consequences rather than method, and should be named without being overweighted. The review was reported widely and in compressed form, and some coverage carried it as a finding that antidepressants do not work. A paper is not answerable for every headline written about it. What belongs in an answer is the clinical consequence, which arrived as patients asking whether they should stop their medication.

PITFALL

Grading a critique by its volume rather than by what it establishes. The loudest objection here is the straw-man one and it is the weakest: it concedes the review's finding and disputes only its novelty. The quietest is the construct objection in the next section, and it is the strongest: it explains why a null result was likely whichever way the biology turns out.

9.6 What the measures could not test, and what was left standing

The strongest move available is not to pick a side. It is to notice that several strands could not have answered the question well even in principle, which weakens the positive literature and the negative conclusion drawn from it in the same movement. Three construct problems recur, and are worth naming in these words.

Concentration is not function: a metabolite concentration in plasma, urine or cerebrospinal fluid is a distant and heavily confounded proxy for serotonin transmission at a synapse in a specific circuit, so a null on the proxy is weak evidence about the thing. Static is not dynamic: the amount of a transmitter present at rest is a different quantity from **release capacity**, the amount a system can release when it is driven, and imaging methods able to probe release with a pharmacological challenge came into use largely after most of the reviewed literature was collected. A positron emission tomography study using an amphetamine challenge, published after the umbrella review, reported reduced serotonin release in a depressed group. It is one small study, it settles nothing, and it is quoted here only to make the structural point: a null covers the measures a literature had, not the measures that came after it. And group is not individual: a

mean difference absent across a diagnostic category assembled from symptoms does not exclude a serotonergic mechanism operating in some subgroup inside it.

Two things the debate never touched should be stated positively; leaving them out makes a script sound as though it has read only the headline. Selective serotonin reuptake inhibitors act on the serotonin system: that is pharmacology, it is not disputed by anyone in this argument, and it is not an aetiological claim. And the neurobiology of depression is broader than any monoamine, with stress-axis, inflammatory, neuroplastic, circuit-level and glutamatergic accounts all under investigation. The Indian guideline's own account is a downstream plasticity account. It puts the typical delay in clinical response at **2 to 4 weeks** and attributes that delay to the time required for gene-level, synaptic and network-level adaptation, stating that antidepressants act by gradually modulating neural circuits rather than by producing immediate relief through neurotransmitter changes alone.

GOOD TO REMEMBER

The most useful thing this controversy leaves behind is not a verdict about serotonin. It is a habit: before repeating any mechanistic explanation of a treatment, ask what would have to be true for it to be right, and whether anyone has checked. That habit transfers to every other explanation offered in a consultation.

9.7 What changes in the consulting room

Nothing about who is offered treatment changes, and everything about how it is explained does. The chemical-imbalance sentence is short, destigmatising and easy to deliver in a crowded outpatient clinic with the family present, which is why it became the default script. It is also a causal claim that the evidence does not support as stated, and it is fragile: a patient who was given it, and who later meets a headline reporting that serotonin was never the problem, has been handed a reason to stop.

The replacement is not "we do not know". Three sentences are defensible and cover what a patient is actually asking. This medicine acts on the serotonin system, which is a statement about the drug rather than about the illness. Why that produces improvement is not fully established, and the current account is about gradual adaptation in brain circuits rather than the topping up of a missing chemical. And the reason to expect benefit is that the medicine has been compared against placebo in randomised trials, which is a different kind of evidence from any theory of cause. Add the expected time to effect and the common adverse effects, and it still fits inside a consultation.

IN INDIA

The 2025 Indian Psychiatric Society depression guideline directs that before prescribing, clinicians provide psychoeducation covering expected benefits, possible risks, anticipated time to effect, and the addressing of myths or misconceptions about antidepressant use, and in its section on suicidal behaviour it directs that the patient and, with their consent, family members be educated about the bio-psycho-social model. Neither instruction endorses a chemical-imbalance script and the second is explicitly broader than any neurochemical account. Note also what is not there. A whitespace-insensitive text search of the four 2025 Indian Psychiatric Society guidelines held for these notes returns no occurrence of "chemical imbalance", "serotonin hypothesis", "serotonin theory" or "monoamine hypothesis" in any of them, while "serotonin" itself occurs 17 times in the depression guideline alone. That is a search of these copies rather than a statement about Indian practice: no Indian guideline position on this controversy, and no Indian survey of what patients are told about the mechanism, is held by this account, and neither should be invented in an answer.

9.8 Answering the question, whatever form it takes

Reduced to five moves, the topic answers as an essay, a short note or a viva.

- Name the design and its inheritance: an umbrella review of systematic reviews, quality-rated, its reach limited by what it collected.
- Give the six strands with a finding attached to each, and mark the sixth as pointing the other way.
- State the reported finding flatly, then attribute every broader conclusion to the authors by name.
- Grade the critiques instead of listing them, and say which would survive a result in the opposite direction.
- Close on the boundary: this is evidence about cause, the efficacy question is answered elsewhere, and what changes is the explanation given to the patient.

MNEMONIC

Design, Data, Authors, Critics, Boundary

Five steps in one direction, and the order is the argument. Design constrains what the data can show; the data are separable from what the authors made of them; the critics attack the middle three and not the first tier; and the boundary is where the marks are, because it is the step everybody skips.

The review reported that a literature does not support a claim; everything stronger than that belongs to somebody, and an answer that says whose is the answer that scores.

10 · Psychedelic-assisted therapy (psilocybin, MDMA) and regulatory status

The examinable object here is not a molecule, it is a package: psilocybin-assisted therapy and MDMA-assisted therapy. Both are a drug given inside a structured psychological intervention with four components. Half the marks sit on reading the evidence honestly: expectancy, functional unblinding, therapist effects, and who was screened out. The other half sit on regulatory status, where several arrangements that read like approval are not one, and only a marketing authorisation licenses a prescription. The pharmacology of these substances as agents of intoxication and misuse belongs to the Substance use disorders notes; this item covers the therapeutic trial and the regulatory status of both agents in three named jurisdictions.

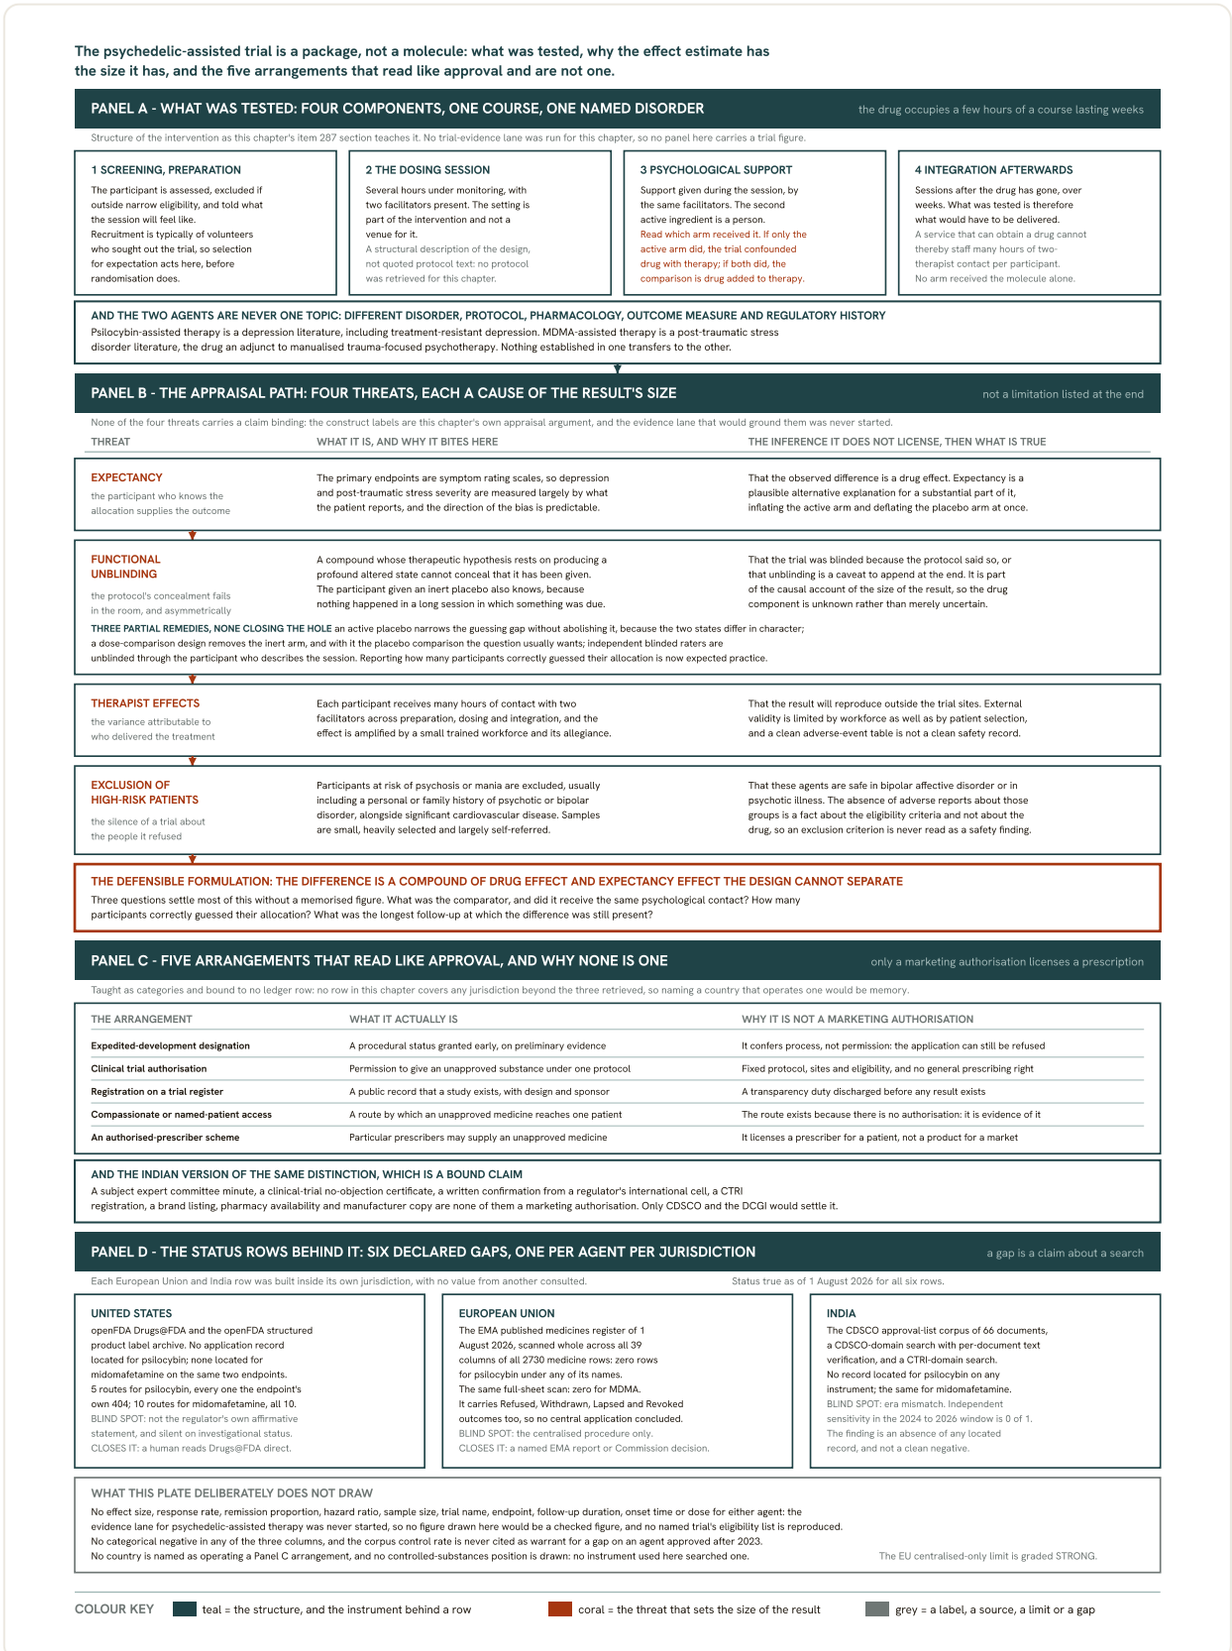


Fig 38.6 Appraising a psychedelic-assisted trial: the package tested, the four threats that set the result's size, and the arrangements that are not a marketing authorisation. The unit of evidence is the package, never the molecule. The four threats are causes of the effect estimate, not caveats after it, so an exclusion criterion is never a safety finding. No effect size, sample size, trial name or follow-up figure is drawn: that lane was never started. Six declared gaps, each naming its instrument: openFDA Drugs@FDA and the label archive, the EMA register scanned whole, and the CDSCO approval-list corpus with CDSCO-domain and CTRI-domain searches, true as of 1 August 2026. (Day 38 status ledger, psilocybin and midomafetamine; the appraisal structure is this chapter's own.)

10.1 What is actually being tested, and why the answer is never "the drug"

Almost every error here is an error of abbreviation. The **intervention package** has four components: screening and preparation, in which the participant is assessed, excluded if outside narrow eligibility, and told what the session will feel like; the dosing session, several hours with monitoring and two facilitators; the psychological support given during it; and integration afterwards. The drug occupies a few hours of a course lasting weeks.

Two consequences follow. The first is attributional: no arm received the molecule alone, so the result cannot be attributed to it, and the intervention's name is psilocybin-assisted therapy, not psilocybin. The second is transfer: what was tested is what would have to be delivered, and a service that can obtain a drug cannot thereby staff many hours of two-therapist contact per participant.

The two agents should never be taught as one topic. The psilocybin literature is a depression literature, including treatment-resistant depression. The MDMA literature is a post-traumatic stress disorder literature, the drug an adjunct to manualised trauma-focused psychotherapy. Different disorder, protocol, pharmacology, outcome measure and regulatory history. Nothing established in one transfers to the other, and a stem that mixes them is testing exactly that.

RULE **Name the package, not the molecule, and never merge the two agents.** The unit of evidence is drug plus preparation plus supported session plus integration, in one named disorder, against one named comparator, over one stated follow-up period. Any sentence that drops a component, or carries a psilocybin depression finding to MDMA in post-traumatic stress disorder, has changed the thing being claimed.

10.2 Two mechanisms, and the reason the trial cannot be blinded

Psilocybin is a prodrug, dephosphorylated to psilocin, an agonist at the **5-HT2A** receptor; the acute state is a marked alteration of perception, affect and self-experience lasting several hours. MDMA is a substituted amphetamine acting principally as a monoamine releaser, with prominent serotonin release, producing emotional openness, reduced defensiveness and sympathomimetic arousal. Its rationale is not that it treats the disorder but that it makes trauma-focused psychotherapy tolerable for people who cannot otherwise engage with the memory.

Read those descriptions as design constraints. A compound whose therapeutic hypothesis rests on producing a profound altered state cannot conceal from the participant that it has been given, so the trial is blinded on paper and unblinded in the room, and every downstream problem descends from it.

These are not agents given daily to steady state. The exposure is one or two supervised sessions, and the benefit is claimed to outlast the drug by weeks or months. That puts a heavy burden on follow-up duration: a result at four weeks says nothing about six months.

10.3 Expectancy and functional unblinding, and what they do to a subjective endpoint

Blinding is not a formality here; it is the load-bearing wall, and it is cracked. **Functional unblinding** means the allocation concealment specified in the protocol fails in practice because the treatment announces itself. It fails almost completely, and asymmetrically: the participant

given an inert placebo also knows, because nothing happened during a long session in which something was expected.

The primary endpoints are symptom rating scales: depression and post-traumatic stress severity are measured largely by what the patient reports. So the participant who knows their allocation also supplies the outcome measurement. **Expectancy** is not a small additive nuisance there; it is a plausible alternative explanation for a substantial part of the observed difference, and its direction is predictable, inflating the active arm and deflating the placebo arm at once.

Three partial remedies exist and none closes the hole. An **active placebo**, producing noticeable but non-psychedelic effects, narrows the guessing gap without abolishing it, because the character of the two states differs. A dose-comparison design, every arm receiving the drug at a different dose, removes the inert arm and the sharpest expectancy contrast, but also the placebo comparison the question usually wants. Independent blinded raters improve on self-report, but a rater interviewing a participant describing their session is unblinded through the participant. Reporting how many participants correctly guessed allocation is now expected practice.

TRAP **Treating unblinding as a limitation to be listed at the end.** It is not a caveat appended to the result; it is part of the causal account of the result's size. The defensible formulation: the between-group difference is a compound of drug effect and expectancy effect the design cannot separate, so the size of the drug component is unknown rather than merely uncertain.

10.4 The therapist is inside the intervention, not outside it

The second active ingredient is a person: each participant receives many hours of contact with two facilitators across preparation, dosing and integration, emotionally charged throughout. **Therapist effects**, the variance in outcome attributable to who delivered the treatment rather than what was delivered, are well recognised in psychotherapy research and amplified here by the small number of trained facilitators, their allegiance to the model, and the intensity of the dyad during an altered state.

Three examinable consequences follow. First, attribution again: if only the active arm receives many hours of skilled psychological contact the trial has confounded drug with therapy; if the placebo arm receives the same contact the comparison is drug added to therapy, not therapy itself. Read which the trial did. Second, replication: delivery depending heavily on individual facilitators is hard to reproduce outside the trial sites, so external validity is limited by workforce as well as patient selection. Third, adverse-event ascertainment: silence collected by an instrument that cannot see a harm is not evidence that the harm did not occur.

PITFALL

Reading a clean adverse-event table as a clean safety record. The suggestibility and dependence on the dyad the model relies on therapeutically are the same features that create the risk of boundary violation, so mechanism and characteristic harm share a root. A safety discussion that lists nausea, headache, transient blood pressure rise and anxiety has described the physiology and omitted the setting.

10.5 Who was enrolled, who was screened out, and what the result therefore covers

Samples here are small and heavily selected. Recruitment is typically of volunteers who sought out the trial, which selects for expectation before randomisation acts. Screening is intensive. And, as a standing feature of the field, participants at risk of psychosis or mania are excluded, usually including a personal or family history of psychotic or bipolar disorder, alongside significant cardiovascular disease.

These trials are silent about the people they excluded. There is therefore no evidence from this literature that psychedelic-assisted therapy is safe in bipolar affective disorder or in psychotic illness, and the absence of adverse reports about those groups is a fact about the eligibility criteria, not about the drug. A candidate who writes that these agents "appear safe, with no cases of induced psychosis reported" has read an exclusion criterion as a safety finding. The same logic bounds every claim: the result describes the population enrolled, the protocol used, the comparator used and the follow-up point reported, and nothing else.

This account prints no effect size, response rate or remission proportion for either agent, and the omission is deliberate. The trial-evidence retrieval for this chapter was not completed, so no figure here would be a checked figure, and a remembered number in this field is characteristically right about the wrong thing: right for a different trial, comparator, endpoint or week. Attach a number only from the paper in front of you, with its population, comparator and time point.

GOOD TO REMEMBER

Three questions settle most viva exchanges here, none needing a memorised figure. What was the comparator, and did it receive the same psychological contact? How many participants correctly guessed their allocation? What was the longest follow-up at which the difference was still present?

10.6 The status ledger for psilocybin and MDMA, jurisdiction by jurisdiction

Regulatory status is recorded on four fields, and a status sentence exists only where a named regulator's record was retrieved: **jurisdiction**, the **exact indication**, the **authoritative source**, and the date the status was true. For both agents, in all three jurisdictions examined, the retrieval found no record. That is six declared gaps, and the correct sentence is not "these agents are not approved" but what was searched, on whose authority a claim would rest, and when.

AGENT AND JURISDICTION	WHAT WAS SEARCHED	WHAT WAS FOUND	STATUS TRUE AS OF
Psilocybin, United States	openFDA Drugs@FDA and the openFDA structured product label archive, 5 query routes, each positive-controlled in the same run	No application record for generic name psilocybin; every route returned the endpoint's own HTTP 404, zero transport failures	1 August 2026
Midomafetamine (MDMA), United States	The same two openFDA endpoints, **10 query routes**, probing midomafetamine, midomafetamine hydrochloride, the full chemical name and the brand-name field	No application record for generic name midomafetamine; all 10 returned HTTP 404, zero transport failures	1 August 2026
Psilocybin, European Union	The EMA published medicines register generated 1 August 2026, scanned whole across all 39 columns of all 2730 medicine rows	Zero rows under any of its names; the register carries Authorised, Refused, Application withdrawn, Withdrawn, Lapsed, Expired, Revoked, Suspended and Opinion outcomes, so this is not merely "not authorised"	1 August 2026
Midomafetamine (MDMA), European Union	The same full-sheet scan of the same register	Zero rows under any of its names; no EU centralised procedure of any kind is on record for this active substance	1 August 2026
Psilocybin, India	The CDSCO published approval-list corpus of 66 documents, a CDSCO-domain search with per-document text verification, and a CTRI-domain search with a positive control	Absent from all 66 lists; the CDSCO-domain search returned zero results at exit status 0, a genuine empty result set rather than a failed query; no indexed CTRI trial record	1 August 2026
Midomafetamine (MDMA), India	The same three Indian instruments	The same finding as psilocybin: no CDSCO record of any kind located on any instrument, and no indexed CTRI trial record	1 August 2026

Two features of that table matter. First, each row was built inside one jurisdiction only: absence in the European Union column was established from EMA sources alone and absence in the India column from Indian sources alone, with no value from another jurisdiction consulted, because a

row may be authored only in the jurisdiction it was grounded in, a declared gap is authorable as a declared gap, and an inferred figure is not authorable at all. Second, the instruments were proved working before their silence was believed: four of four hosts round-tripped a known-present record before any miss was read as a fact, with certificate verification never disabled. An earlier run reported nothing anywhere for any agent; the cause was a broken certificate store, not a fact about a drug.

2730	66	10	0
EU REGISTER ROWS SCANNED	CDSCO APPROVAL DOCUMENTS READ	OPENFDA ROUTES, MIDOMAFETAMINE	RECORDS LOCATED, ALL JURISDICTIONS

Each instrument has a measured blind spot. The American search covers openFDA Drugs@FDA and the label archive only; it is not the regulator's own affirmative statement that no application exists, and it is silent on investigational and breakthrough designations, which appear in neither instrument. That silence is the trap: the same null comes back whatever stage a programme has reached, so it is no evidence that nothing is under way. What either instrument indexes, and what it leaves out, is a property of the instrument, and this chapter did not search it: what was measured is the scope sentence above and nothing wider. A claim about the coverage of either endpoint would rest on the regulator's own description of it, retrieved and dated, and none was retrieved here. Closing the row: a human must read Drugs@FDA directly and, if there is no application, say so on the regulator's authority rather than on this probe's silence.

The European search covers the centralised procedure. A medicine authorised nationally through the decentralised or mutual-recognition route would not appear in the register at all, so absence here means no EU centralised procedure of any kind on record, which is narrower and weaker than "not available anywhere in the European Union". That limit was graded: STRONG for both, because for a small molecule outside the mandatory centralised scope a national authorisation is genuinely possible and wholly unmeasured here. A second weakness: the confirming second instrument was applied to 8 of 8 grounded rows and 0 of 8 gap rows, so it fired only where the first instrument had already found something, and no European gap in this chapter has independent corroboration. A claim here would have to rest on an EMA European public assessment report or a European Commission implementing decision, named as such; the phrase "approved internationally" may never be substituted for a named regulator.

The Indian instrument is the weakest of the three, and its weakness is measured, which is why the India cells are gaps and not negatives. Of 20 drugs taken as approved and marketed in India, 17 appear in the corpus and 3 do not: agomelatine, molnupiravir and trospium. More decisive is the era mismatch. Every one of the 17 present controls appears in a document dated 2023 or earlier, independent sensitivity in the 2024 to 2026 window is 0 of 1, and no sentence may cite the 17 of 20 figure as warrant for a gap on any agent approved after 2023. The corpus also contradicts itself: CDSCO's own subject expert committee minutes of 16 April 2025 state that another agent in this chapter, whose status row is another item's, was already approved in the country, yet that agent appears in none of the 66 lists. Hence the corrected wording on both India rows: "clean negative" overstates what an instrument with a measured 15 per cent miss rate, blind to three of

its own positive controls, can deliver, and the finding is an absence of any located record rather than a clean negative.

TRAP **Writing "psilocybin is not approved in India".** That claims knowledge of the regulator's whole record, and this instrument cannot support it. What it supports is narrower and checkable: no CDSCO approval record for psilocybin was located in the 66 published approval lists, in a CDSCO-domain search that returned a genuine empty result set, or in a CTRI-domain search, as of 1 August 2026. The general form: name the register, name the date, name what would close the row, and let the absence be an absence.

The falsification pass that tested this chapter's status rows against their primary documents scored all six not applicable, on the stated ground that a declared gap carries no indication text, so an instrument that compares text against a primary document has nothing to test: a gap is not a claim about a product, it is a claim about a search. They have been checked for what they say about the search, never independently confirmed as absences.

10.7 Five things that are not a marketing approval

Each of the five below is routinely written up, in press and review articles, in language indistinguishable from approval. A marketing authorisation is a decision by a named regulator that a specific product may be placed on the market for a stated indication.

MECHANISM	WHAT IT ACTUALLY IS	WHY IT IS NOT APPROVAL
Expedited-development designation	A procedural status granted early, giving a sponsor closer regulatory interaction and faster review of a programme for a serious condition	Granted on preliminary evidence, it confers process, not permission: the application can still be refused or withdrawn, with no product reaching the market
Clinical trial authorisation or permission	Permission to give an unapproved substance to defined participants under a defined protocol at defined sites	It authorises research, and its conditions are the opposite of a marketing licence: fixed protocol, sites and eligibility, no general prescribing right
Registration on a trial register	A public record that a study exists, with its design, sponsor and eligibility criteria	A transparency requirement discharged before results exist. It says a study was declared, not that it worked or that anything was licensed
Controlled, compassionate or named-patient access	A route by which an unapproved medicine reaches a specific patient with no alternative, usually with individual authorisation and reporting conditions	The route exists precisely because there is no marketing authorisation. It is evidence of the absence it works around
An authorised-prescriber or special-access scheme	An arrangement permitting particular approved prescribers to supply an unapproved medicine to particular patients	It licenses a prescriber for a patient, not a product for a market. The product remains unapproved, and the scheme's paperwork usually says so

Two distinctions belong beside it. First, controlled-substances law, on which this chapter searched nothing. Every instrument behind these notes is a drug-approval or drug-trial instrument, and not one is a statute, a scheduling notice or a controlled-substances register, so nothing retrieved here establishes what any schedule permits or forbids, what altering one would do, or how a schedule stands beside a marketing authorisation in any jurisdiction. A claim about any of that would rest on the statute itself and on the body that administers it, named, quoted and dated. Until one is read, treat the schedule and the licence as two questions put to two records, and answer only the one these notes hold a record for. Second, chapter hygiene: a subject expert committee minute, a clinical-trial no-objection certificate, a written confirmation from a regulator's international cell, a CTRI registration, a brand listing, pharmacy availability and manufacturer copy are none of them a marketing authorisation, and the only bodies whose record would settle the Indian question are CDSCO and the Drugs Controller General of India.

This section names no country as operating any of the five mechanisms, deliberately: no ledger row covers any jurisdiction beyond the three retrieved, so attaching a named scheme to a named country would be writing from memory.

IN INDIA

The position to hold, and it is narrow. As of **1 August 2026**, no CDSCO or DCGI approval record for psilocybin was located in the 66 published approval lists, in a CDSCO-domain search, or in CTRI, and the finding for MDMA is identical. Say that rather than "not approved", and say what would change it: an affirmative CDSCO or DCGI new-drug permission, a current CDSCO-approved label, or an entry in an official approval list. A CTRI registration, if one appears, establishes that a study exists and does not close the approval question. No instrument used for this chapter looked at the controlled-substances position for either agent in India, so nothing here establishes what that position is or what it requires, and the schedule number stays out of your answer until the statute and the body that administers it have been read. And when a patient or family arrives having read about these treatments abroad, name the gap rather than filling it: no located Indian marketing authorisation for either agent, a multi-session therapy package rather than a prescription, and trials that excluded several groups who reach Indian outpatient departments in large numbers.

10.8 What a complete answer contains

Four moves answer this item, in any form.

- Define the intervention as a package, name its four components, and attribute the result to the package, not to the molecule.
- Give functional unblinding and expectancy as the reason the effect estimate has the size it has, not as a limitation appended afterwards, and say what a design can and cannot do.
- Bound the evidence by the population enrolled: small, selected, self-referred, screened to exclude those at risk of psychosis or mania, so no safety claim is available for bipolar or psychotic illness.

- State regulatory status on the four fields, register named, date attached, gap declared, and separate approval from designation, trial permission, registration, access arrangement and prescriber scheme.

MNEMONIC

Jurisdiction, Indication, Source, Date

The four fields of a status sentence, in writing order. A claim missing any one is not weaker, it is different: without jurisdiction a claim about the world, without indication a claim about a molecule, without a source a claim about the writer's memory, without a date a claim about now, forever.

RULE **Absence anywhere in this section is a statement about a search, never proof of absence.** A queried register's null is evidence and worth stating; it is not proof, and the difference has been demonstrated here. One agent in these notes, brexanolone, was recorded as having no United States record until multi-route probing retrieved application NDA211371, ZULRESSO, approved 17 June 2019 and now of Discontinued marketing status: the record carries no openfda block at all, so three of five query routes returned NOT_FOUND while a fourth returned it. Strengthen a null by widening the routes, not by repeating one: an independent recheck of the psilocybin row on a separate run repeated the probe across 9 routes, including free-text and wildcard searches on both endpoints, returned zero on every one, and had its own positive controls return hits in the same run.

The package is what was tested, unblinding is why the number is uncertain, exclusion is what the number does not cover, and a designation, a trial permission and an access scheme are not a licence.

11 · Artificial intelligence and machine learning in psychiatry

A model does not diagnose anybody; it returns a number, and the whole of this topic is what that number is allowed to mean. The marks sit not on architectures but on four questions asked of any reported figure: what was the model predicting, in whom, against what, and had the data ever been seen before. Candidates lose them in three predictable places: an accuracy figure quoted without its population, a good discrimination statistic read as a good decision, and a performance claim treated as a regulatory permission. The statistics themselves, sensitivity, specificity, predictive values and the appraisal of a diagnostic accuracy study, belong to the research methodology and biostatistics notes; what follows is only what changes when the index test is a fitted model rather than an assay.

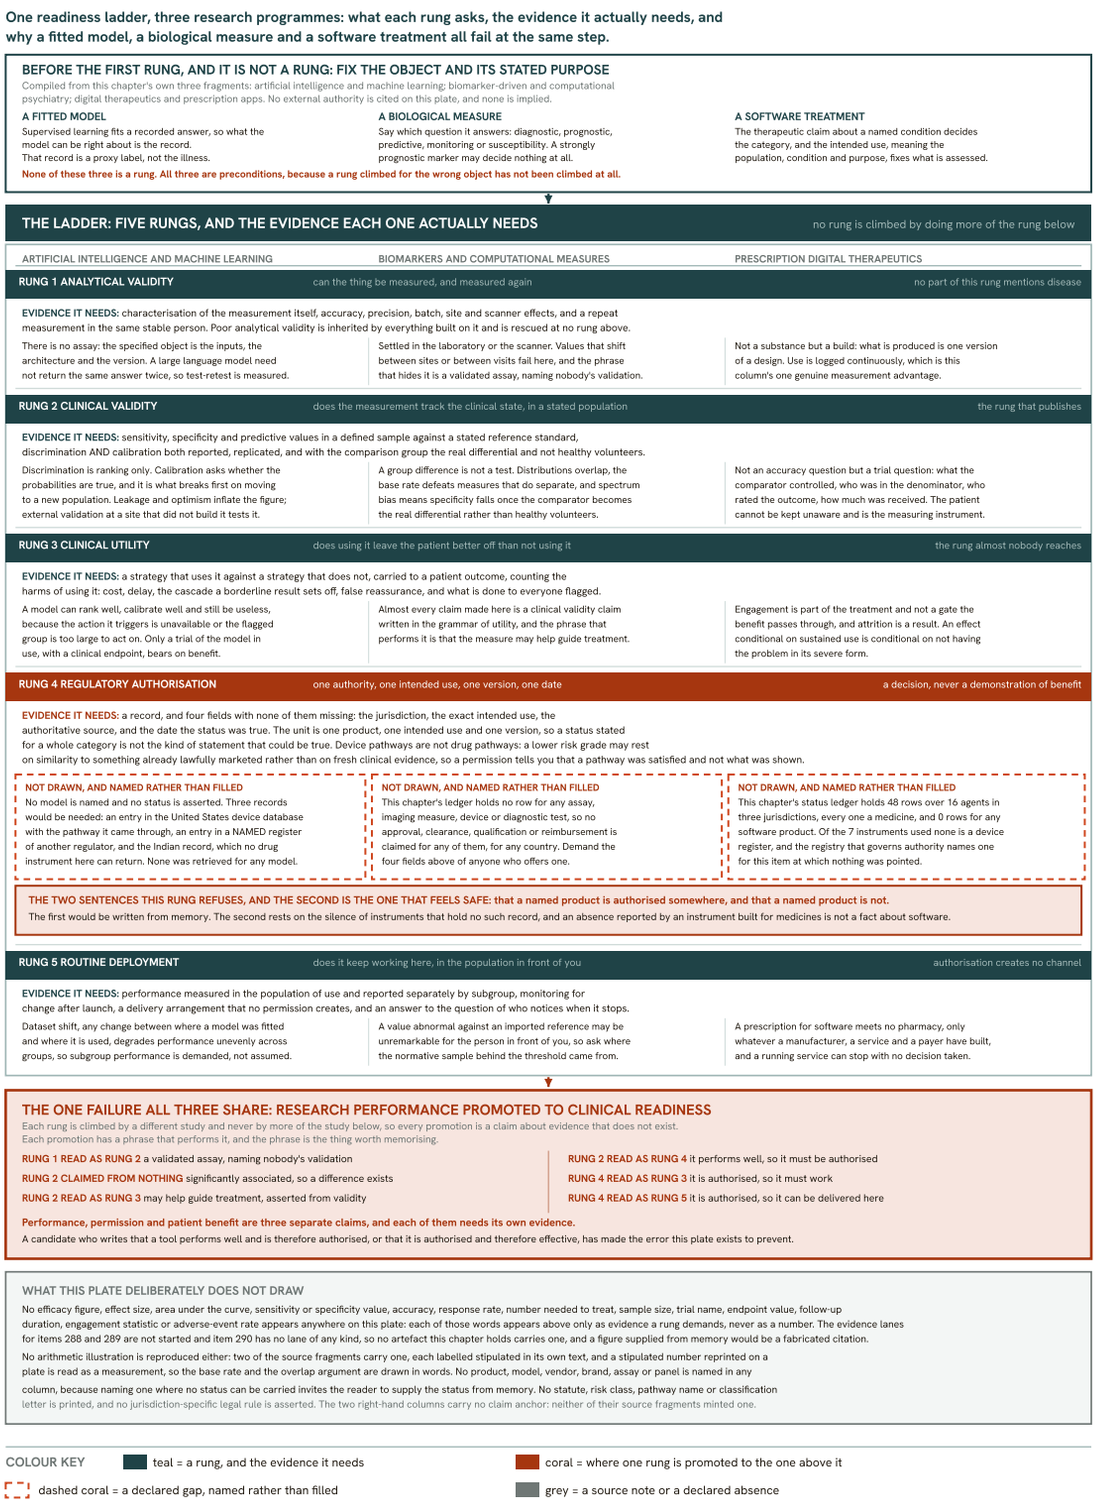


Fig 38.7 One readiness ladder for artificial intelligence, biomarkers and prescription digital therapeutics: analytical validity, clinical validity, clinical utility, regulatory authorisation, routine deployment. Above the first rung the object and its purpose are fixed: supervised learning fits a recorded answer, and that record is a proxy label rather than the illness. Each rung is climbed by a different study, never by more of the study below, so discrimination and calibration settle rung 2 and neither touches rung 3, and external validation at a site that did not build the model answers what leakage and optimism flatter. Rung 4 is drawn EMPTY and named rather than filled: no product is named and no status asserted. Rung 5 is where dataset shift bites. Source: this chapter's own three fragments, items 288, 289 and 290; the evidence lanes for 288 and 289 read as not started and item 290 has no lane at all.

11.1 What the machine is actually predicting

Almost every clinical model here is **supervised learning**: the algorithm is shown examples in which the answer is already recorded, and fits a function from the inputs to that recorded answer. Everything it can ever be right about is therefore the record, not the illness, the first thing to establish about any reported result and the one most often skipped.

The recorded answer is a **proxy label**. In psychiatry the proxies available at scale are administrative or documentary: a diagnostic code entered at discharge, a prescription of a particular class, a rating-scale score above a cut-point, a chart entry flagging risk, a readmission within a stated window. None of these is the construct: a code records what a clinician wrote down in a system that pays for codes. A model trained on codes learns coding practice, which varies by clinician, by unit and by country, and reproduces it faithfully.

The construct validity of the label is a ceiling on everything downstream: no split, no architecture and no sample size repairs a target that does not mean what the paper says it means. And the label determines what the number transfers to, so a model that predicts a recorded diagnosis of depression predicts documentation, and a claim that it predicts depression has quietly changed the target between the methods section and the abstract.

11.2 The six things a performance number carries, or it is not a finding

A performance figure with nothing attached is uninterpretable, not merely imprecise. Six items travel with it, each naming a distinct way a good-looking number goes wrong.

WHAT MUST BE STATED	WHY IT IS NOT OPTIONAL	WHAT A MISSING ENTRY HIDES
The model	Inputs, architecture and version define the thing being reported	A retrained or updated model is a different model with a different number
The target	The exact label, how it was defined and by whom	The proxy problem above, and a target that drifts between sections
The population and setting	Who was included, recruited how, where and when	Case mix and prevalence, which move predictive values without moving the model
The split method	How the evaluated data were separated from the fitted data	Leakage and optimism, the commonest cause of an inflated figure
The comparator	What the model is being judged against	A model beating nothing, or beating a deliberately weakened baseline
External validation status	Whether the number comes from data the model never saw, from a site that did not build it	The single largest source of the gap between published and delivered performance

The comparator row is the one candidates forget and an examiner can mark hardest. A model that predicts an outcome better than chance has been compared with nothing. The comparisons that matter are the ones a service would otherwise use: the existing clinical judgement, an established

structured instrument, or a simple regression on a handful of routinely available variables. That last comparator is unglamorous and frequently wins, which is why a paper omitting it has omitted the informative arm.

RULE **A performance figure belongs to one model, one target, one population, one split, one comparator and one validation status, and it does not travel without all six.** Move any one of them and the number is no longer about the same thing. This is the whole discipline of the topic, and why a sentence of the form "artificial intelligence can detect depression with high accuracy" is not a weak claim but an empty one: it names none of the six.

11.3 Discrimination is not calibration, and neither one is usefulness

Discrimination is the model's ability to rank: given one person who has the outcome and one who does not, how often does it give the first the higher score. That is what the area under the receiver operating characteristic curve measures, and its floor is worth fixing in memory: a model at **0.5** ranks no better than a coin. Discrimination is a property of ordering only, and says nothing about whether the predicted probabilities are true.

Calibration is that second question: among the people the model gave a twenty percent probability, did about twenty percent have the outcome. A model can rank well and be badly calibrated, systematically over-predicting risk in everybody. The distinction has a direct clinical consequence: a threshold is set on the probability, not on the ranking, so a poorly calibrated model puts the operating threshold in the wrong place even when its ranking statistic looks strong. Calibration is also the property most damaged by moving a model to a new population, and the least often reported.

Neither statistic answers the question a service actually has: whether acting on the output leaves patients better off than not acting. That is a question about consequences at a chosen threshold: what is done to the people flagged, what it costs, and what is missed in the people not flagged. Decision-analytic measures of net benefit address it, and no ranking statistic substitutes. A model can discriminate well, calibrate well and still be useless, because the action it triggers is unavailable, or the group flagged is too large to act on.

11.4 The base rate, and why an accurate model flags mostly false alarms

This is the most productive arithmetic in the topic, and where risk-prediction claims in psychiatry come apart. Sensitivity and specificity are properties of the model. Predictive values are not: they depend on how common the outcome is in the people the model is applied to, and psychiatry's most-wanted outcomes are rare. What follows is a worked example with numbers chosen for arithmetic, not measured in any study, and worth reproducing in an answer as a worked example.

Take an outcome occurring in **1 in 100** of those screened, and a model with **80 percent** sensitivity and 80 percent specificity, which would be a respectable psychiatric model. In 10,000 people, 100 have the outcome and the model finds 80 of them. Of the 9,900 who do not, it wrongly flags 1,980. So 2,060 people are flagged, 80 of them with the outcome.

1%	80%	3.9%	99.7%
ASSUMED BASE RATE, NOT MEASURED	ASSUMED SENSITIVITY AND SPECIFICITY	RESULTING POSITIVE PREDICTIVE VALUE	RESULTING NEGATIVE PREDICTIVE VALUE

Read both ends of that strip. About 26 people are flagged for every one who has the outcome, so whatever the flag triggers is delivered, and paid for, overwhelmingly in people who will not have the event. Meanwhile the negative predictive value looks magnificent and is uninformative: with an outcome this rare, predicting no event in everybody already achieves 99 percent, so the model must be judged against that baseline and not against zero. Both figures move whenever the model is applied to a population with a different base rate, which is why a screening claim from a specialist inpatient sample does not survive the move to a general clinic.

PITFALL

Quoting accuracy as the headline figure when the outcome is rare. Accuracy is the proportion of all predictions that were correct: with an outcome occurring in one percent of people, a model predicting no event for everybody is 99 percent accurate and clinically worthless. An answer repeating an accuracy figure on an imbalanced outcome without noticing the imbalance has walked into the oldest trap in the field.

11.5 How the number gets inflated: leakage, optimism and the wrong split

Published performance is usually higher than delivered performance, and the reasons are structural rather than dishonest. **Data leakage** is information about the answer reaching the model through a route that will not exist in use. Its forms are separate examinable errors: preprocessing such as scaling or feature selection performed on the whole dataset before the split, so the held-out data influenced the fitting; a feature that is a consequence of the outcome rather than a predictor of it, such as a treatment started because the event had already occurred; the same patient contributing records to both sides of the split; and a time window that lets the model see data recorded after the moment the prediction is supposed to be made.

Optimism is the second mechanism and it needs no error at all. When many models, many feature sets and many thresholds are tried and the best result is reported, that result is partly a selection from noise, and it shrinks on new data. The same happens when the held-out set is consulted repeatedly during development, at which point it has joined the training process. The problem is worst where candidate features greatly outnumber patients, the ordinary condition of neuroimaging and genomic work in psychiatry.

Both are properties of how the evaluation data were separated, which makes the split method the most informative line in a methods section. The ladder runs weakest to strongest.

SPLIT	WHAT IT HOLDS OUT	WHAT IT STILL CANNOT RULE OUT
Random split of records	Individual rows	The same patient on both sides, and every form of leakage
Patient-level split	Whole patients	Drift over time, and anything peculiar to that one site
Temporal split, trained on earlier data and tested on later	The future of one site	Everything specific to that site: case mix, coding habits, systems
External validation at a site that did not build the model	Site, population and practice	That the model, once acting on care, still performs as measured
Prospective evaluation with the output withheld from clinicians	Retrospective advantage	Whether acting on it changes any patient outcome

Even the bottom rung falls short. A model can be externally validated, prospectively confirmed and well calibrated and have no demonstrated effect on any patient outcome, because performance and benefit are different endpoints and only a trial of the model in use measures the second. Reporting and appraisal frameworks exist and are worth naming: TRIPOD plus AI for the reporting of prediction model studies, PROBAST plus AI for risk-of-bias appraisal, SPIRIT plus AI and CONSORT plus AI for protocols and reports of trials of these models, and DECIDE plus AI for early live clinical evaluation. The mark goes to knowing what each covers, since a paper satisfying the reporting standard is completely reported and not thereby correct.

TRAP **Reading an internal cross-validation figure as though it were an external result.** Cross-validation resamples inside one dataset, so it estimates how the fitting procedure behaves there and nothing else. It does not test a different site, a different case mix, a different coding culture or a different era, and where preprocessing or feature selection was done before the folds were drawn it does not even estimate that honestly. Examinably: internal validation bounds optimism within one sample, and only evaluation on data from an institution that had no part in building the model bears on whether it transfers.

11.6 Large language models: a different tool with a different failure surface

A **large language model** is not a clinical classifier and should not be appraised as one. It is trained to produce plausible continuations of text, so its outputs are fluent by construction and true only incidentally, and the failure modes follow from that rather than from any defect in a particular product. It produces confident, well-formed statements that are false, including references that do not exist. It tends to agree with the framing it is given, so a leading clinical question returns a supporting answer. It emits no calibrated probability alongside its answer, removing exactly the quantity the previous sections were about. And it need not return the same answer twice, so test-retest reliability cannot be assumed and has to be measured.

Two further points are examinable because counter-intuitive. Performance on written examination questions is a poor guide to clinical performance, partly because a question with five options and one defensible answer is not the task of the clinic, and partly because material that has circulated on the internet may sit in the training data, so a benchmark score can reflect

exposure rather than reasoning. And a general-purpose model given clinical text is not thereby a medical device or a validated instrument; what it is depends on the claim made for it and how it is deployed.

GOOD TO REMEMBER

The nearest-term uses in psychiatry are administrative rather than diagnostic: drafting documentation from a recorded consultation, summarising a long record, coding, translation support. These carry real benefit through time returned to the clinician, and a distinct risk profile, since a fabricated or subtly altered detail in a summary enters the record and is read as fact downstream. The safeguard is structural rather than technological. Output is a draft; the clinician who signs it owns every word in it, and the review has to be real, because a summary that is right nine times out of ten trains the reviewer to stop reading it.

11.7 Regulatory status: what this section refuses to say, and why

Nothing in this fragment states that any named model, product or system is approved, cleared, licensed or authorised anywhere. That is a deliberate refusal and the correct examination position, because the approval status of software is per product, per intended use and per version, and there is no class-wide approval of artificial intelligence in medicine of which any individual tool is an instance. The status ledger built for this chapter covers medicines only, and a drug ledger licenses no statement about a device.

Three records would have to be held before any regulatory sentence could be written about a named model, and naming them demonstrates the point better than reciting a status. The first is an entry in the United States device database for that specific product with the pathway it came through, since clearance against a predicate device, authorisation as a novel low-risk device and full premarket approval are different standards with different evidence behind them. The second is an entry in a named register of another regulator, identified rather than assumed from the first, because a marking obtained in one jurisdiction is not a permission in another. The third is the Indian record, and it is where an answer should be most careful, because the instruments this chapter searched are drug instruments: nothing retrieved here establishes which Indian pathway governs a software product, which body administers it, or what class and licence any given product carries. Those are matters of record, and this chapter holds no such record and did not look for one. No such record was retrieved for this chapter for any model, so no product is named here and no status is asserted for any.

Two errors of reasoning persist even when the records exist. A regulatory permission is evidence that a standard was met for a stated intended use; it is not evidence that patients do better, and the evidence supporting a clearance is often accuracy against a reference standard rather than any outcome. And because the object approved is a version, a model that continues to learn, or is retrained, raises the question whether the thing in use is still the thing assessed, which is why change-control for adaptive software is a live regulatory problem rather than a settled one.

TRAP **Sliding from a performance claim to a permission, or from a permission to a benefit.** These are three separate objects and each transition is a separate error. A published accuracy figure is a measurement in a sample; a regulatory clearance is a decision about one product for one stated use; a demonstrated patient benefit needs a prospective study of the model in use with a clinical endpoint. A candidate who writes that a tool is approved and therefore effective, or that it performs well and is therefore approved, has made the error the whole topic exists to prevent.

11.8 Bias, the Indian question, and what a complete answer contains

A model cannot be called fair, and the reason is that fairness is not one quantity. Several formal definitions of fairness exist and conflict mathematically, so a model satisfying one will generally violate another, and a claim of fairness without a stated definition and a measurement is not a finding. The mechanisms producing unequal performance are, however, straightforward. Groups under-represented in the training data are fitted less well. The label itself is socially patterned, since who receives a recorded diagnosis, who is flagged as a risk and who is admitted all depend on who is doing the recording. And **dataset shift**, meaning any change between the setting where the model was fitted and the one where it is used, degrades performance in ways that fall unevenly across groups. The remedy in an answer is not an assurance but a demand: performance reported separately by subgroup, in the population of use.

IN INDIA

The motivation is real and the strongest argument in the topic: the specialist workforce is far smaller than the need, and any tool that extends a clinician's reach is worth taking seriously. The caution is equally concrete. Almost all of this literature is built on structured electronic records from North American and European health systems, and there is no reason to expect a model fitted to those records to hold here, where much of the record is on paper or free text, where clinical documentation is often multilingual and code-switched, where case mix and pathways to care differ, and where the base rates in any given clinic differ from those in the development sample. Nothing about Indian performance is asserted anywhere in this section, because no Indian external validation of any named psychiatric model was retrieved for this chapter. Two positions hold regardless of the evidence. Accountability for a clinical decision rests with the registered practitioner and is not transferred to a vendor by the use of a tool. And identifiable patient information entered into a general-purpose service leaves the clinic, so the confidentiality obligation has to be settled before the tool is used and not after; which statutory instrument governs that transfer is a legal question this chapter has not established and therefore does not answer.

The topic reduces to four decisions.

- Establish the target before the number: what was recorded, by whom, and whether it is the construct or a proxy for it.
- Ask what the number was measured against, both in comparator and in split, and treat internal validation as an estimate of optimism rather than of transfer.
- Take the base rate into the arithmetic before accepting a screening claim, because predictive values belong to the population and not to the model.

- Keep performance, permission and patient benefit as three separate claims, each needing its own evidence.

MNEMONIC

Target, Population, Split, Comparator

The four questions that dismantle any reported figure, and the order to ask them in: what was actually being predicted, in whom, how were the evaluated data kept away from the fitting, and against what was the result judged. A claim that survives all four may then be asked the fifth and hardest: whether acting on it helps anyone.

A model's number is a measurement in one sample of one recorded label; it becomes a clinical claim only when the target, the population, the split, the comparator and the validation status are all named with it.

12 · Brain-based, biomarker-driven and computational psychiatry approaches

Psychiatry has published thousands of candidate biomarkers and holds none of them as a diagnostic test for any of its own disorders. Explaining that sentence is the item, not listing findings: grading any biological claim, from a scan to a polygenic score to a fitted model, on three questions asked in order. Can the thing be measured, does the measurement track the clinical state, and does using it leave the patient better off. Collapsing them into one is the characteristic error of the area, made as often dismissively as enthusiastically. Two neighbours are bounded out: appraising an algorithm, its dataset and its external validation belongs with the artificial intelligence and machine learning material in this chapter, and measures that have entered diagnostic use in neurodegenerative illness belong with the Dementias and neurocognitive disorders notes.

12.1 What a biomarker claim is actually claiming

A **biomarker** is a characteristic measured objectively as an indicator of a biological process or of a response to an intervention. Measured objectively excludes the rating scale, which reports experience, so a depression severity score is not a biomarker of depression however well it performs. Indicator excludes mechanism: a marker need not cause anything, and most psychiatric candidates are downstream or non-specific.

The claim must state which question the measure answers. A diagnostic marker says whether the person has the condition now, needing a reference standard and a stated comparison group. A prognostic marker says what course the person will take whatever the treatment. A **predictive** marker says which of two treatments will suit this person better, an interaction between marker and treatment. A monitoring marker says whether the state has changed, needing sensitivity to change rather than group separation. A susceptibility marker says how likely onset is in someone currently well, a statement about a population, not a person's future.

That pair is the distinction most reliably examined and most reliably muddled. A marker can be strongly prognostic and useless for choosing: knowing a patient will do badly whatever you do changes nothing. Showing a marker is predictive needs patients on different treatments compared across levels of the marker, a far more demanding design. A complete claim carries five things: the measure, the question, the population, the comparator, and the decision it feeds. Supplying the missing ones from expectation invents them.

12.2 The spine: analytical validity, clinical validity, clinical utility

Analytical validity asks whether the measurement is a good measurement, settled in the laboratory or the scanner by the properties in the table below. No part of it mentions disease, and poor analytical validity cannot be rescued downstream: the noise is inherited by every analysis built on the measure.

Clinical validity asks whether the measurement relates to the clinical state in a defined population. **Discrimination** and **calibration** are the components routinely dropped. Discrimination is the ability to rank a person with the condition above one without it, summarised by the area under a receiver operating characteristic curve. Calibration is whether a predicted probability matches observed frequency, so a one in ten risk holds about one time in ten. A model can discriminate respectably and be badly calibrated, and a badly calibrated risk is unusable at the bedside while reading well in the paper.

Clinical utility asks the only question a patient cares about: does using this test lead to better outcomes than not using it. It compares a strategy including the test against a strategy without it, carried to a patient outcome, counting the harms of testing: cost, delay, the cascade a borderline result sets off, false reassurance, and the effect of being told you carry a biological abnormality. Almost every claim made for a psychiatric biomarker is a statement about clinical validity, sometimes only analytical validity, written in the grammar of clinical utility.

QUESTION	WHAT IT ASKS	WHAT ESTABLISHES IT	HOW IT FAILS, AND THE PHRASE THAT HIDES IT
Analytical validity	Does the assay or scan measure the quantity accurately and reproducibly	Laboratory characterisation: accuracy, precision, detection limit, batch, site and scanner effects, repeat measurement	Values that shift between sites or between visits. Hidden by "a validated assay", which names nobody's validation and no population
Clinical validity	Does the measurement track the clinical state in a stated population	Sensitivity, specificity, predictive values, discrimination and calibration, in a defined sample, against a stated reference standard, replicated	Separation that shrinks once the comparison group resembles the real differential. Hidden by "significantly associated", which says a difference exists, not that a test works
Clinical utility	Does using the test improve what happens to the patient	A strategy with the test against a strategy without it, carried to a patient outcome, counting the harms of testing	No such study exists for the marker in question. Hidden by "may help guide treatment", asserting the third question from evidence for the second

RULE **Analytical validity, clinical validity and clinical utility are asked in that order, and none implies the next.** A perfectly reproducible assay may bear no relation to any clinical state, and a measure that separates two groups beyond doubt may change no decision and improve no outcome. Grade every claim by naming which of the three it has established, and say plainly that the others are unaddressed rather than assuming they follow.

12.3 Why a group difference is not a test

A group difference says one distribution's average sits away from another's; a test has to place an individual on one side of a line. A difference can be beyond statistical doubt while the distributions overlap almost completely, because a large enough sample makes a small separation reliably detectable. Reliably detectable and usefully separated are different properties, and the first is what most published biomarker findings have.

The second defeater is the base rate, and it defeats measures that do separate. Take a test right in 90 percent of people who have a condition and in **90 percent** of those who do not, applied where **1 person in 100** has it. Of 10,000 tested, 90 of the 100 with the condition test positive, and 990 of the 9,900 without it also test positive. A positive result is therefore far more often wrong than right: **fewer than 1 in 10** positives is a true positive. The prevalence did the damage, not the test, and no gain in sensitivity undoes it, because the false positives come from the enormous group without the condition.

90%	90%	1%	8%
SENSITIVITY, STIPULATED	SPECIFICITY, STIPULATED	HOW COMMON THE CONDITION IS	CHANCE A POSITIVE IS TRUE

Those four figures are a worked illustration, not a measurement of any real test. The shape is what to carry: at low prevalence positive predictive value collapses while negative predictive value becomes excellent, so a marker can be useful for ruling out and useless for ruling in.

The third defeater is **spectrum bias**. Performance is first measured in clear, established, often severe illness against healthy volunteers; the clinic asks the hard question, an early, partial, comorbid presentation against the other diagnoses on the differential. Move the comparison group to the real differential and specificity falls, sometimes to nothing, because the marker was never specific to the disorder, only to being unwell. An accuracy figure must travel with the comparison group it was earned against; without it the figure is uninterpretable rather than merely imprecise.

TRAP **Repeating a research accuracy figure as though it described the clinic.** Published accuracy is typically obtained in a balanced sample against healthy participants, in people whose diagnosis is already certain. The clinic has none of those properties: low prevalence, a comparison group that is the differential, and patients presenting early and ambiguously, which is why the test was wanted. The defensible sentence names the sample and the comparator beside the number, and states that performance in an unselected population is unknown unless measured there.

12.4 The reference standard problem, and why a research framework is not a diagnosis

The deepest difficulty is not the measure, it is the yardstick. A diagnostic biomarker is validated against a **reference standard**, and in psychiatry that standard is a clinical diagnosis made by interview against the current manuals. Two consequences follow.

The first is a ceiling: a measure judged against clinical diagnosis can at best reproduce it, its apparent accuracy bounded by the reliability of the standard rather than the biology, so where two experienced clinicians would disagree the marker is scored against a coin. The second is a filter running the wrong way: a measurement identifying a real biological grouping that cut across the manual's categories would be scored inaccurate and abandoned, because the criterion for success is agreement with the classification it contradicts. That circularity is why agreement with diagnosis is a weak result and predicting what it does not, such as course or treatment response, a strong one.

Dimensional research frameworks, which organise study around functional constructs measured across levels from genes and circuits to behaviour and self-report, and the hierarchical taxonomies derived from the covariance of symptoms, exist because of that problem: they free research from the assumption that the manual's categories are natural kinds. They are research frameworks, carrying no diagnostic threshold, no validated decision rule and no treatment recommendation attached to any dimension, and none allocates treatment better than a diagnosis made by interview.

GOOD TO REMEMBER

Phenomenology is not a placeholder waiting for biology to replace it. It remains the measurement instrument in psychiatry and the standard every candidate marker is validated against, so it cannot be displaced without a study showing better patient outcomes rather than better agreement. The viva position is neither that biology will soon supply the diagnosis nor that it offers nothing: these approaches produce constructs, stratification hypotheses and research designs, and no claim to have replaced the interview has been made good for any primary psychiatric disorder.

12.5 What is replicated, what it licenses, and what it does not

Group-level findings here are real; the discipline is to say what each licenses. Every row below has been reproduced, none is a test, and the fourth column carries the marks.

MEASURE FAMILY	WHAT REPLICATES AT GROUP LEVEL	WHAT THAT LICENSES	WHAT IT DOES NOT LICENSE
Structural imaging	Enlarged lateral ventricles in chronic schizophrenia, among the field's oldest and most reproduced findings, and small average reductions in cortical and subcortical volumes across several disorders	A statement about averages, and a research target	Individual inference. Distributions overlap heavily with comparison groups and the differences are not specific to one diagnosis
Functional imaging and connectivity	Group differences in task activation and in resting connectivity across mood, psychotic and anxiety disorders	Hypotheses about circuits, and stratification questions worth testing prospectively	A scan for a patient. Effects are small against within-person and between-scanner variability
Electrophysiology	Reduced auditory event-related potential amplitudes and reduced mismatch negativity in schizophrenia, among the most replicated findings in psychiatry	Research into information processing, and endophenotype work in relatives	Diagnosis or exclusion in an individual. The reductions are neither specific nor separating
Peripheral and inflammatory measures	Modest average elevations of circulating inflammatory markers in a subgroup of people with depression	A stratification hypothesis, to be tested rather than applied	Diagnosis, and any causal statement. Confounded by adiposity, smoking, infection and physical comorbidity, and specific to no psychiatric disorder
Neuroendocrine	Group differences in cortisol regulation, the historical example being the dexamethasone suppression test proposed for melancholic depression	A permanent teaching case: it separated groups, and specificity collapsed once the comparison group was other psychiatric patients	Diagnostic use, which was proposed, tried and abandoned
Common-variant genomics	Polygenic scores associated with case status for schizophrenia and for bipolar disorder at the group level, robustly and repeatedly	Research into shared architecture across diagnoses	Individual prediction, diagnosis or counselling. Variance explained is small, distributions overlap almost completely, transfer across ancestries is poor

The one place where fluid and imaging measures have crossed into diagnostic use is neurodegenerative illness, taught with the Dementias and neurocognitive disorders notes; nothing about their performance transfers to a primary psychiatric disorder, and citing them as

proof that psychiatry now has biomarkers is a category error. And this section makes no regulatory statement of any kind: the chapter's status ledger holds one row per drug per jurisdiction and no row for any assay, imaging measure, device or diagnostic test, so no approval, clearance, qualification or reimbursement claim appears here for any country. Demand of anyone who offers one the four fields the chapter uses throughout: the jurisdiction, the exact intended use, the authoritative source, and the date the status was true. An intended use is bound to the jurisdiction that granted it, and a test available in one country establishes nothing about another.

12.6 Computational psychiatry: what a model parameter is, and what it is not

Computational psychiatry states a hypothesis about a mechanism as a formal model, fits it to an individual's behaviour or brain data, and treats the fitted quantities as candidate measures: how a person processes information, in numbers comparable across people.

Two strands run under the name. The theory-driven strand builds a **generative model**, equations describing how choices or percepts are produced, and estimates its parameters for each person: reinforcement learning models yield learning rates and sensitivity to reward and punishment; evidence accumulation models yield an accumulation rate and a decision threshold, separating a slow response caused by cautious responding from one caused by inefficient processing; Bayesian and predictive coding models yield the weight given to prior expectation against incoming evidence, describing hallucination and delusion as an altered balance rather than an unexplained content. The data-driven strand predicts without committing to a mechanism, and its discipline is taught with the artificial intelligence and machine learning material.

A fitted parameter is an estimate produced by fitting a model you chose to data you collected, a property of the model and the data together rather than a substance inside the patient, so changing the model changes it. Four requirements follow beyond the ordinary ones. **Parameter recovery** comes first: simulate data from known values, refit, and check that the known values are recovered, because a parameter unrecoverable from data the model itself generated cannot be interpreted in a patient. Identifiability next: two mechanisms producing indistinguishable behaviour cannot be told apart by fitting. Third, model comparison is relative, selecting the best of the models written down, and the best of a bad set is still that. Fourth, prediction must be out of sample, with any selection of features or participants kept inside the training data.

Apply the spine and a computational measure reads cleanly. Analytical validity becomes parameter recovery plus the stability of the estimate when the same person repeats the task. Clinical validity becomes whether the parameter separates clinical groups or predicts an outcome. Clinical utility becomes whether allocating treatment by the parameter beats allocating it clinically. Computational psychiatry is substantial at the first two and effectively silent at the third, and why the third is a different study and not a bigger one is the examinable point.

12.7 Reliability, replication, and the claim that a biological measure is unbiased

Two mechanical facts explain most of the disappointment here. The first is a reliability trap in task design: a task is optimised to produce a robust average difference between conditions, done

by making everyone respond as similarly as possible, deliberately minimising the between-person variation an individual measure needs. So a group effect that replicates without fail can sit beside a between-session correlation close to nothing in the same individuals, both consequences of one design choice rather than a contradiction. Because reliability caps the correlation a measure can show with anything else, an unreliable measure cannot correlate strongly with symptoms or outcome however real the relationship. A measure built for group inference must be rebuilt before it serves as an individual marker.

The second is that discovery is not validation. Clustering a sample into subtypes always returns clusters, whose number and shape depend on the measures entered, the preprocessing and the algorithm. Several prominent psychiatric subtyping solutions have not been recovered in independent samples. A subtype derived in one dataset is a hypothesis until an independent dataset recovers it and it predicts something outside the data that generated it.

PITFALL

Treating poor test to retest reliability as a minor technical caveat. It is not a footnote to the finding but a limit on it. If a measure taken twice in the same stable person gives substantially different answers, it cannot support an individual clinical decision whatever its group statistics look like. Ask what the between-session agreement was before asking what the between-group difference was, and note how often that question is unanswered.

The claim that measuring biology removes human bias is false, and the four routes by which bias enters are a genuine discriminator. The label carries it first: a marker trained against clinician-assigned diagnoses learns whatever systematic patterns those clinicians produced, reproducing the bias with a more authoritative surface. The sample carries it second: performance is a property of the population studied, shown most clearly by polygenic scores, whose performance falls substantially when transferred from the ancestry group they were derived in. The instrument third: scanner, batch and site effects are real, and where sites differ in whom they recruit, a site effect and a demographic effect cannot be separated afterwards. Access fourth: who is tested determines who can benefit, and private testing is distributed by income rather than need.

IN INDIA

The transfer problem is not abstract here. Reference intervals, normative neuropsychological data and polygenic scores are largely derived from populations in which South Asians are thinly represented, so a value abnormal against an imported reference may be unremarkable for the person in front of you, and a risk score performs worst precisely where it is newest to the market. Two consequences. A group frequency is a prior about a population and never a result about a patient, so no measure may be read off ancestry. And before treating a value as abnormal, ask where the normative sample behind that threshold came from; where local normative data exist, use them.

12.8 Where a biological test does change management, and what a complete answer contains

An answer finishing on scepticism has missed the clinical half of the item. The tests that earn their place are not tests for psychiatric disorders but for conditions that both imitate a psychiatric presentation and carry a different treatment, and their utility is easy to demonstrate because a positive result changes management at once.

The categories follow from the principle rather than a memorised panel: endocrine disturbance, particularly thyroid disease; nutritional deficiency; infective and inflammatory causes; antibody-mediated encephalitis, characteristically a rapidly evolving first episode of psychosis with atypical features, seizures, movement abnormalities or autonomic instability; inherited metabolic disease such as Wilson disease in a young patient with psychiatric and movement symptoms together; structural and epileptic causes; and substance and medication effects, commoner than all the rest. Justify each test by the decision it could change, and notice that this is the one part of the topic where the third question of the spine is genuinely satisfied.

IN INDIA

Two situations recur in outpatient practice. A family arrives having paid privately for a brain mapping report, a neurotransmitter panel or a genetic wellness profile and asks what it means. The question is not whether the number sits outside its range but which decision would change if it were different, and where the honest answer is none, saying so is the correct clinical act, with an explanation of why the test could never answer the question it was sold as answering. An unvalidated test is not neutral: it has cost the family money and it generates false reassurance. The second situation is the opposite failure, the treatable mimic missed because nothing was investigated.

TRAP **Writing that a scan or a blood test can now diagnose a psychiatric disorder, or that it soon will.** Neither half is defensible. There is no validated transdiagnostic blood test or brain scan for a primary psychiatric disorder; a group-level association does not determine an individual diagnosis; and a research classification organises study rather than replacing the interview. The forecast is not safer than the claim, because a prediction about validation is a claim about evidence that does not exist. Say what has been established, name which of the three questions it answers, and state that the others are open.

Four moves answer this item in whatever form it is asked.

- Grade the claim on the spine: analytical validity, clinical validity, clinical utility, named in order, with the one actually established stated and the others called open.
- Give the two arithmetic reasons a group difference is not a test, overlapping distributions and the base rate, with spectrum bias as the reason published accuracy falls in clinic.
- Name the reference standard problem: a marker validated against clinical diagnosis is capped by it and discarded if it disagrees, so predicting course or treatment response is the stronger design.
- Treat a computational parameter as a property of the model and the data, needing recovery, identifiability, out-of-sample testing and test to retest reliability before it measures anything

in a person.

MNEMONIC

Measure, Separate, Decide, Benefit

The first three are the spine and the fourth is the study almost nobody has done: does the result change a decision, and does that change leave the patient better off. Every biomarker claim sits at the step where its evidence stops.

A biomarker is graded on three questions asked in order, a group difference is not a test, and the yardstick every candidate is judged against is still the clinical interview.

13 · Digital therapeutics and prescription apps in psychiatry

A prescription digital therapeutic is a treatment whose active ingredient is software, and almost every difficulty follows from that sentence. The clinical question is how an intervention delivered on a patient's own device is evaluated, when its effect depends on how much of it the patient does, when no convincing dummy version can be built, and when the user cannot be kept unaware of what they are using. The regulatory question is what kind of object it is, since it is not a medicine and does not reach a market by a medicine's route. Marks are lost in three places: appraising the product as though it were a drug, reading an effect measured in the patients who kept using it as the effect of the product, and supplying a regulatory status from memory because none was written down. Two neighbours share this chapter's readiness ladder and neither is rebuilt here: appraising a fitted model, its dataset and its external validation belongs with the artificial intelligence and machine learning material, and grading a measure on analytical validity, clinical validity and clinical utility belongs with the biomarker material.

13.1 What the object is, and the four things it is not

A **prescription digital therapeutic** is a software product that makes a therapeutic claim about a named condition, delivers the intervention to the patient rather than to the clinician, and is gated behind a clinician's authorisation. The claim decides the category, not the content: two products can ship identical exercises and be different objects entirely, because one asserts treatment of a named condition and the other does not.

Four neighbours are routinely collapsed into it. A wellness or self-help application may carry the same material, makes no therapeutic claim, and is therefore assessed by nobody. A telemedicine platform delivers a clinician at a distance, so the treatment remains the clinician's and the software is a channel. A decision-support tool acts on the clinician's judgment and the patient never touches it, which is the object the artificial intelligence and machine learning material appraises. Software supplied alongside a medicine carries no therapeutic claim of its own, because the claim belongs to the drug.

The sentence that does for software what a molecule does for a drug is its **intended use**: the population, the condition and the clinical purpose, written down. Everything is assessed against

it, and a product whose intended use is amended is a different product.

13.2 Why it is a device and not a drug, and the six things that follow

A medicine is a substance and its manufacture reproduces that substance. Software has no substance: what is assessed is a design together with an intended use, and what is produced is a build. Hence the category **software as a medical device**, software intended for a medical purpose that achieves it by processing information rather than by chemical or physical action on the body.

First, the unit of authorisation is one product, one intended use and one version, so there is no class of digital therapeutics of which a product is an instance, and a class-level statement of status is not the kind of statement that could be true. Second, risk classification drives the evidence burden: device frameworks grade by the harm a failure could cause, and lower grades may rest on similarity to something already lawfully marketed rather than on a fresh trial, so two products may be marketed on entirely different evidence. A permission tells you a pathway was satisfied, not what was shown. Third, there is no generic: a second manufacturer cannot show its software equivalent to another company's, so it builds a different product with its own authorisation and evidence.

Fourth, **change control** is continuous rather than exceptional: a tablet's formulation is frozen at approval, software ships updates, and since the object assessed was a version, the question after launch is whether the thing in the patient's hand is still the thing assessed. Fifth, withdrawal has no drug analogue, because a dispensed medicine stays with the patient, whereas a therapeutic depending on a running service can stop working and the patient loses the treatment with no clinical decision taken. Sixth, being authorised creates no way to deliver it: a prescription for a molecule meets a pharmacy, while a prescription for software meets whatever arrangement a manufacturer, a service and a payer have built, and where none exists it cannot be filled.

RULE **The thing authorised is one product, for one intended use, at one version; move any of the three and the permission does not travel with you.** Three errors share that shape: a claim about a category, a claim about a wider population than the intended use names, and a claim about the current release when the evidence came from an earlier one, the last invisible unless somebody asks which version was studied.

13.3 What this chapter can see, and what it structurally cannot

The most useful thing this section can give a candidate is a worked example of an absence that means nothing. This chapter carries a status ledger of one row per agent per jurisdiction, each row carrying the same four mandatory fields, the jurisdiction, the indication or a declared gap in its place, the authoritative source or that same gap, and the date the status was true. It holds **48** rows over **16** agents in three jurisdictions, every one a medicine. The number of rows describing any software product is **0**.

That is not an oversight, and reading more of the same documents would not close it. The instruments recorded as used number **7**. Five can see nothing but medicines: two interfaces onto a national drug application and labelling archive, a register of centrally authorised medicines

with its assessment pages, and a corpus of published new-drug approval lists. The sixth is a whole-domain search of the drug regulator's own website, pointed in this build only at drug names. The seventh searches the national trials registry, which records research and never marketing status. A drug application archive cannot return a device record, because that is not the kind of record it holds. The five archives are silent about software by construction, the other two by what they were pointed at.

Two facts in this chapter's working files make that sharper. The source registry that governs which document is authoritative names, for artificial intelligence and digital therapeutics specifically, a device database, and that database appears neither among the instruments recorded as used nor among those reachable but not yet run. And the grounding lanes name the drug items and do not include this item at all, not even a status lane. The absence here is not a search that came back empty; it is a question none of the records these instruments hold could answer.

48	16	0	7
ROWS IN THE CHAPTER STATUS LEDGER	AGENTS COVERED, ALL MEDICINES	ROWS FOR ANY SOFTWARE PRODUCT	INSTRUMENTS USED, NO DEVICE REGISTER

The falsification pass built to test those rows compares a stored indication text against a cached primary drug document, and a software product would not fail those instruments: having no drug label to read, it is inapplicable to them. An instrument that cannot fail is not a test.

So no product is named in this section, and no approval, clearance, authorisation, licence, marking, date or jurisdiction is asserted for anything. Two sentences are forbidden and the second is the one that feels safe: that a named product is approved somewhere, and that a named product is not approved somewhere. The first would be written from memory; the second would rest an assertion on the silence of instruments that hold no such record.

TRAP **Filling the regulatory blank yourself, in either direction.** A reader who meets a section naming no product and no status will supply both, because a refusal reads as an omission. Half-remembered regulator statements are the commonest fabrication here, and the mirror error is worse, concluding from silence that nothing has been authorised anywhere. The examinable move is to name the record that would settle it, state that such a record is per product, per intended use and per version, and say none was retrieved.

13.4 How the evidence differs from a drug's evidence

The trial machinery looks familiar and diverges at six points.

QUESTION	A MEDICINE	A PRESCRIPTION DIGITAL THERAPEUTIC
What is administered	A fixed quantity of a substance	An opportunity to do a course of work, which the patient may not take
How exposure is known	Inferred from prescriptions and self-report	Logged continuously, module by module
What varies between patients	Absorption, metabolism, adherence	How much was received at all, from none to all
What blinding is possible	Allocation and the treatment itself	Allocation only; the patient knows what they are using
Who measures the outcome	Often a rater unaware of allocation	Usually the unblinded patient, on the same device
After authorisation	A frozen formulation	Successive versions, of which the studied one may no longer be in use

The exposure row cuts both ways. Precise logging is a real advantage, since treatment received is measured rather than assumed. The same precision creates the field's most attractive error: because exposure is known and varies enormously, an analysis restricted to those who completed the programme is available, tempting and severely biased, since completers are the engaged and engagement predicts outcome for reasons unconnected to the software. The randomised comparison lives in the analysis that keeps everybody allocated to it.

13.5 Engagement and attrition, which do something a pill's adherence does not

A tablet delivers its benefit on being swallowed. A therapeutic delivers its benefit only when the patient does the work inside it, so the active ingredient is the patient's own activity. **Engagement** is therefore not adherence renamed: adherence is a gate the benefit passes through, whereas engagement is part of the benefit, which is why an effect size is uninterpretable without a statement of how much of the intervention was taken up.

Attrition then does four separable things. It empties the denominator, steeply, since disengagement is heaviest early. It is not random, so where the arms lose different sorts of patient the randomisation that made the comparison fair has quietly been undone. It is itself an outcome rather than a nuisance, because a treatment patients stop using is not working. And it is entangled with the state the product exists to change, since low motivation, poor concentration and hopelessness are features of the conditions treated and also the commonest reasons a person stops opening an application. An effect conditional on sustained use is conditional on the patient not having the problem in its severe form.

Two further effects have no drug counterpart. Novelty decays, so an effect measured across a short window need not describe the same product used for a year. And the human contact around a trial is a co-intervention: visits, reminders and a coach who telephones when use lapses deliver

support, so a therapeutic evaluated with that scaffolding is a different intervention from the software used alone, and the unsupported version is usually what a service plans to deploy.

GOOD TO REMEMBER

Prescribing one of these is prescribing a course of work, so time the review to where the work stops, not to the end of the programme. Silent abandonment is the ordinary outcome, and a patient who has quietly stopped is often assumed by everyone to be in treatment. Ask what the patient did, in what week they last did it, and what got in the way, and treat disengagement as clinical information, not non-compliance.

13.6 Why the control condition is unusually hard to build

Every comparator here fails to control something, and naming what each leaves uncontrolled is the appraisal skill worth carrying.

COMPARATOR	WHAT IT CONTROLS	WHAT IT LEAVES UNCONTROLLED
Waiting list	Time and spontaneous change	Attention, expectancy, structured activity and being monitored, so the difference is an upper bound, not an estimate
Usual care	Whatever the service delivers	What that is: unstandardised, differing between sites, and in an under-resourced service possibly very little
Attention-matched or sham application	Screen time, contact, expectancy, daily use	Whether it is truly inert, since anything engaging enough to be used is doing something, and whether patients believed it
The same treatment delivered by a person	The content, isolating the mode of delivery	Nothing important, which is why it is the informative arm, the most expensive and much the rarest

The **digital placebo** names the problem: the benefit that comes from using something at all, from the daily ritual, the attention and the expectation of improvement. A **sham comparator** must satisfy three requirements in tension: plausible enough to be believed, engaging enough not to be abandoned faster than the active arm, and inert. Anything meeting the first two risks failing the third, and differential dropout between a dull sham and an interesting therapeutic reintroduces the bias the sham was built to remove.

Blinding compounds it. Allocation can be concealed at randomisation, but a participant cannot be kept unaware of what they are using for weeks, and the outcome is usually their own self-report on the same device, so the person who knows which arm they are in is also the measuring instrument. That is a structural weakness rather than a defect of any one study, and a larger sample does not repair it. Two consequences: an outcome assessed by someone unaware of allocation, or one not self-reported at all, carries much more weight here than a self-rated scale;

and framing is part of the intervention, because content presented as prescribed and authorised generates an expectancy the same content offered as a self-help download does not.

TRAP **Reading a waiting-list-controlled result in completers as the effect of the product.** It stacks two errors working the same way. The waiting-list comparison leaves attention, expectancy and structured activity uncontrolled, so it measures the difference between doing something and doing nothing rather than the difference this product makes; restricting the analysis to completers then removes everyone in whom the product failed to hold attention. Ask what the comparator was, who was in the denominator and who assessed the outcome before accepting any figure.

13.7 What an Indian service would have to establish, and what a complete answer contains

The lens here is the product: what the evidence shows, what kind of regulated object it is, and what a prescription for it would mean. Access, equity, service integration, adoption and programme monitoring for digital mental health in this country are the subject of the Community psychiatry and public health notes, because whether a product works and whether a system can deliver it are answered from different evidence.

IN INDIA

Six things must be established before a service deploys one, and none may be inferred from an application store listing, a company's own description, or a permission granted in another country. First, what kind of regulated object the product is here and on whose authority, with the document named and dated, since no instrument in this chapter's set can return that record and nothing in this section supplies it. Second, whether the evidence was generated in a population and language resembling the patient, and whether the content was rebuilt for that language or merely translated, since a therapeutic is made of words in a way a tablet is not. Third, which version the patient will receive, and whether it is the version studied. Fourth, who holds the data it generates and where that goes, since a therapeutic collects a clinical record continuously as a condition of working. Fifth, what happens when the patient stops, and who notices. Sixth, whether the human support that surrounded the product in its trials can be supplied here, because without it the intervention deployed is not the one evaluated.

Reduced to four decisions, the item is answerable in any form.

- Identify the object first: a therapeutic claim about a named condition, delivered to the patient, gated by a clinician, assessed as one version of one product.
- Appraise it as a device trial: what the comparator controlled, who was in the denominator, who rated the outcome, how much was received.
- Treat engagement as part of the treatment and attrition as a result, and refuse any effect size quoted without both.
- Keep evidence, permission and deliverability separate, and where no regulatory record has been retrieved, name the missing record instead of guessing what it would say.

MNEMONIC

Claim, Version, Comparator, Completion

Claim is what makes software a regulated object and fixes what may be said about it. Version is what was assessed, and the question nobody asks. Comparator is what the trial controlled, and so what its number means. Completion is how many were still using it when the outcome was measured, both a result and the limit on every other result.

The active ingredient is the patient's own work, the object authorised is one version of one product for one intended use, and an absence reported by an instrument built for medicines is not a fact about software.

Individual newer agents

14 · Asenapine: pharmacology and clinical use

One active substance, two routes, and four regulatory records that do not agree on what the drug is for. Asenapine is the most compact item in this chapter and the easiest to answer wrongly with confidence: a candidate who knows it from the bipolar algorithm writes an indication correct in one country and false one border away. What follows is the whole-agent profile: the formulations, the indication set jurisdiction by jurisdiction, and the discipline that keeps the two fastened together. Receptor pharmacology, the sublingual dosing detail and asenapine's place in the bipolar phase algorithm are taught in the Mood disorders: pharmacotherapy notes.

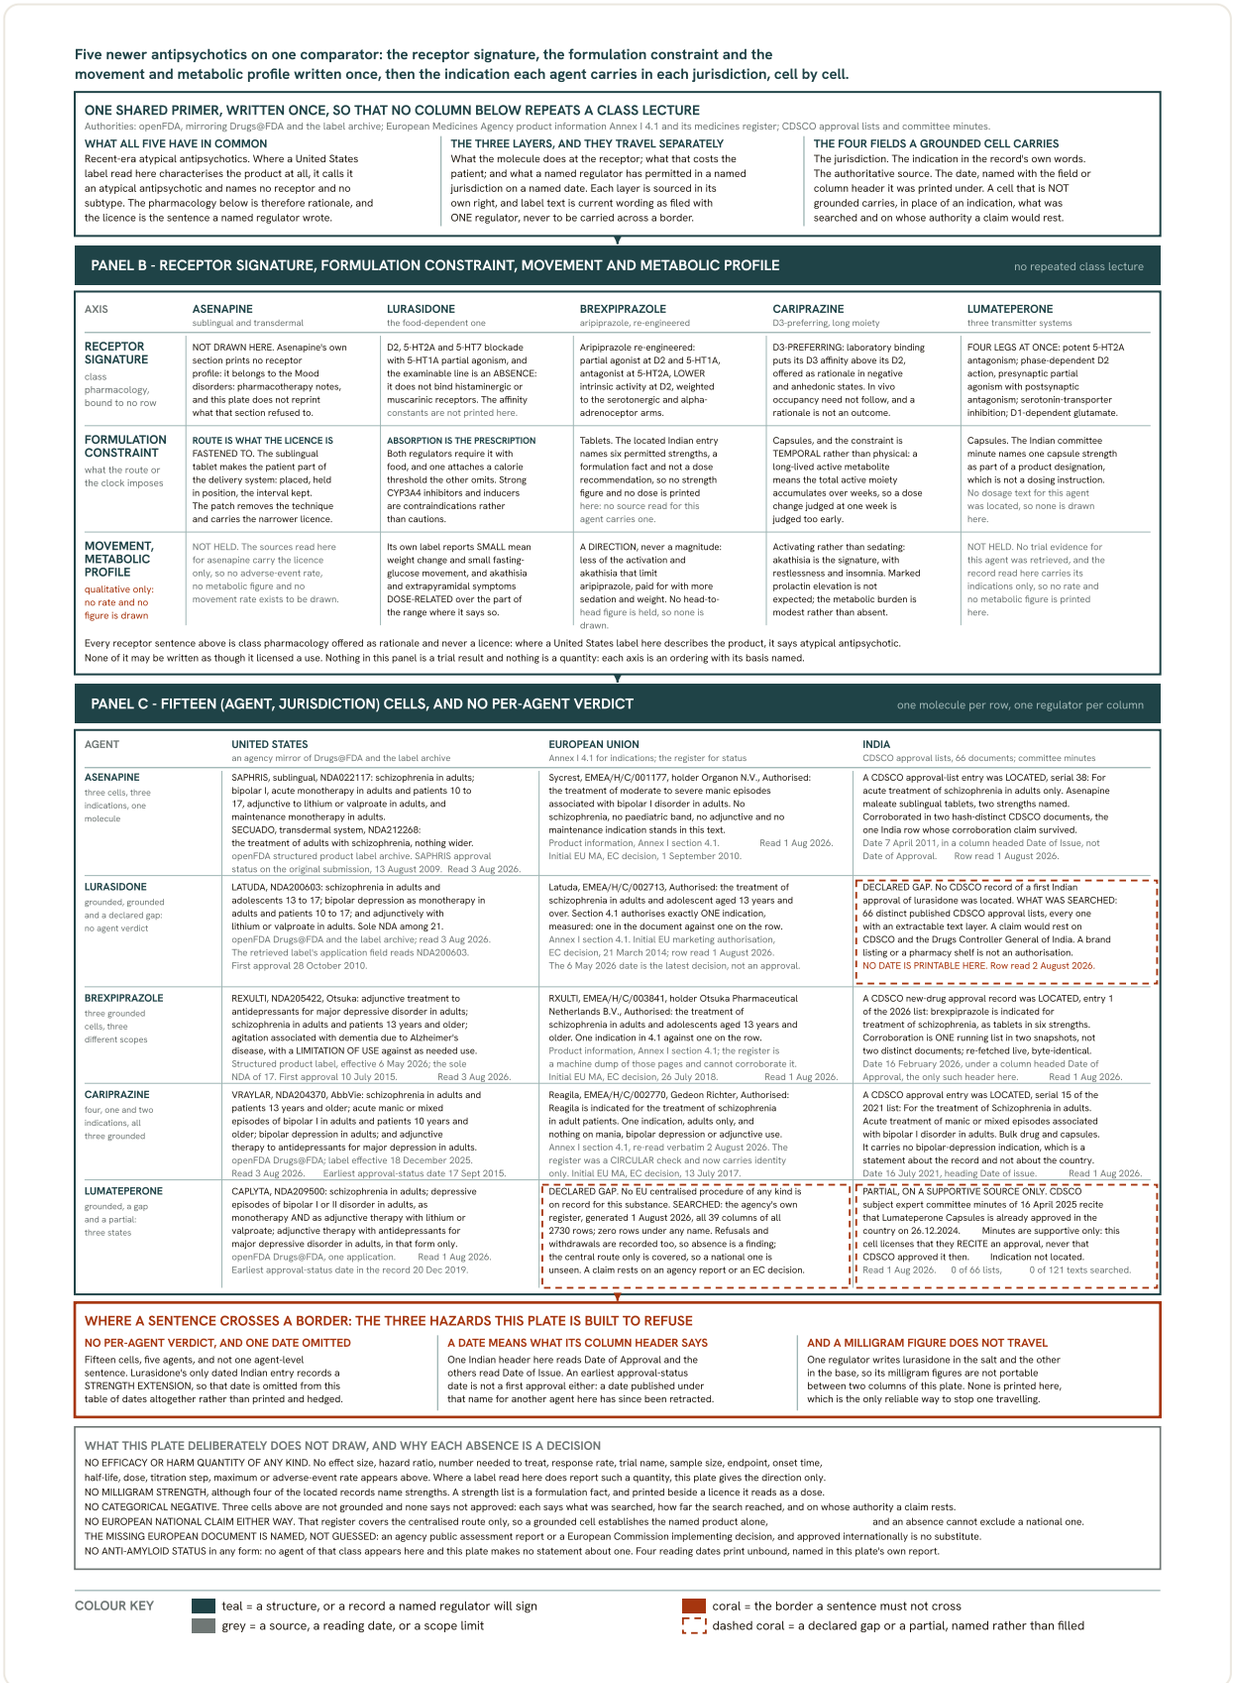


Fig 38.8 **Five newer antipsychotics on one comparator: fifteen (agent, jurisdiction) cells, not five verdicts.** Where a United States label read here describes the product it says only atypical antipsychotic, so the receptor row is rationale and the licence is the regulator's sentence. Movement and metabolic effects are qualitative: no rate, dose or onset time appears. The states differ inside one molecule: lurasidone's Indian record is a declared gap, lumateperone's recites an approval it does not establish, and no cell is a per-agent verdict. A cell that is not grounded names what was searched. (openFDA; European Medicines Agency Annex I 4.1 and medicines register; CDSCO lists and committee minutes.)

14.1 Two American applications, and why they are not one drug

The United States holds two applications for asenapine, and they are not two names for one product. The **innovator application** is a **sublingual tablet**, SAPHRIS: application NDA022117, sponsor ABBVIE, whose original submission carries approval status dated **13 August 2009**, in three strengths, 2.5 mg, 5 mg and 10 mg expressed as asenapine base. The second is a **transdermal system**, SECUADO: application NDA212268, sponsor HISAMITSU, whose original submission carries approval status dated **11 October 2019**, as three patches delivering 3.8 mg, 5.7 mg and 7.6 mg over 24 hours.

They do not license the same thing. The sublingual label reads: SAPHRIS is indicated for schizophrenia in adults, and for bipolar I disorder as acute monotherapy of manic or mixed episodes in adults and pediatric patients 10 to 17 years of age, as adjunctive treatment to lithium or valproate in adults, and as maintenance monotherapy treatment in adults. The transdermal label reads, whole: SECUADO is indicated for the treatment of adults with schizophrenia. One route carries schizophrenia and a three-part bipolar I indication reaching to age 10; the other adult schizophrenia alone.

Three constructs inside that indication are routinely collapsed together. **Acute monotherapy** of manic or mixed episodes is licensed in adults and in patients aged 10 to 17. **Adjunctive treatment** to lithium or valproate is licensed in adults only, and so is **maintenance monotherapy**. The paediatric band sits on one of the three and not on the other two, so an answer that calls the drug licensed in adolescents without saying for what has widened a licence by omission.

RULE **A regulatory sentence names four things or it is not one: the jurisdiction, the exact indication in the record's own words, the authoritative source, and the date the status was true.** On this agent the fourth field alone splits: the American row was read on 3 August 2026, the European and Indian rows on 1 August 2026. Copying one reading date onto the others is the same error as copying an indication across a border.

14.2 The row that blended three products, and how it was caught

The most transferable thing about this item is not the drug; it is the way its record went wrong. The first American row was built by querying the regulator's interface on the generic name, returning five applications, three of them abbreviated applications for generic products, and a set spanning SAPHRIS, SECUADO and generics. A falsification pass rejected it on two instruments. The earliest date across a mixed set is not a first approval date but the earliest among whatever the query returned; and a set spanning more than one product means the label and dates retrieved may describe different products.

What the repair found is worse than a wrong date. The row cited SECUADO, carried the date belonging to SAPHRIS, and published the sublingual bipolar text that SECUADO does not carry: citation, date and indication text each pointed at a different application. Nothing in its appearance signalled that. It now holds the innovator with its own matching date. The repair instruction is the method: query per application rather than per generic name, paginate to

exhaustion, and write one row per application with its own sponsor, brand, approval date and label, never merging applications into an agent-level row.

The row was then tested rather than trusted. Three text instruments were run against the primary label: that the row text is verbatim, that it is not the primary text cut short, and that the count of indications carried matches the document. All three passed, at 680 characters against 680 and one indication against one. It also carries a caution about itself: the interface is an agency-operated mirror of the approval database and the label archive rather than the application itself, and the wording is current label text as filed, which must never be carried across a jurisdiction.

PITFALL

Searching a regulator by drug name and reading the first date that comes back. A generic-name query returns innovator and generic applications together, and the earliest date in that heap belongs to whichever record happened to be earliest, not to the drug. For asenapine the heap held three generic applications out of five. Ask which application, then which product, then which date, in that order.

14.3 Route is what the indication is attached to

Asenapine is given sublingually because a swallowed tablet does not deliver a useful amount of drug; the pharmacokinetic reason, the dose and the timed post-dose instruction belong to the Mood disorders: pharmacotherapy notes. A sublingual tablet makes the patient part of the delivery system: the tablet has to be placed, left in position, and the interval afterwards kept, or the dose administered is not the dose prescribed. A transdermal system moves that burden onto an application site.

That is why a transdermal asenapine exists, and why the trade is not free. The route that removes the technique is the route with the narrower licence: adults with schizophrenia and nothing more, against a sublingual label that also carries bipolar I disorder. A patient with bipolar I mania who cannot manage a sublingual tablet cannot simply be moved to the patch.

One boundary must be stated. No trial-level evidence is held for this agent in this build: item 291 carries a status lane and no evidence lane, so nothing here names a pivotal trial, a sample size, an endpoint, an effect size or an adverse-event rate. The comparison above is of licences and not of efficacy.

14.4 The European Union: one product, one indication, not the American one

Europe holds one centrally authorised asenapine product: Sycrest, product number EMEA/H/C/001177, status Authorised, whose **marketing authorisation holder** is Organon N.V., with an initial European Union marketing authorisation, by European Commission decision, dated **1 September 2010** and a latest European Commission decision of 25 February 2026.

The indication, verbatim: Sycrest is indicated for the treatment of moderate to severe manic episodes associated with bipolar I disorder in adults. Read what is absent. No schizophrenia. No

paediatric age band. No adjunctive indication with lithium or valproate, and no maintenance indication. And a severity qualifier, moderate to severe, that the American text does not carry.

Where that sentence comes from is examinable in itself. Its authority is the **Summary of Product Characteristics**, Annex I section 4.1. The row originally took its indication from the agency's published medicines register, citing the assessment-report page as an independent second instrument; the register is a machine dump generated from those same pages, so the corroboration was circular and blind to the very failure it existed to catch. The register and that page now corroborate product identity and authorisation status only. When the eight grounded European rows here were re-read against the product characteristics, six differed and two were already verbatim; asenapine was one of the two. A row that turns out to be right is not a row that has been checked.

One limit travels with every European sentence here: this is an authorisation for the named product through the **centralised procedure**. It says nothing about other products containing the same active substance that may be authorised nationally through the decentralised or mutual-recognition procedures, which the register does not cover by construction. The defensible sentence is about Sycrest, not about asenapine everywhere in the Union.

TRAP **Writing the American bipolar indication into a European or an Indian answer.** The European authorisation carries moderate to severe manic episodes associated with bipolar I disorder in adults and nothing else. The Indian entry carries acute schizophrenia in adults and nothing else. Neither carries the American bipolar text. Each record is correct where it was issued and wrong where it was not, and the wrongness is invisible unless the regulator is named.

14.5 India: the one corroboration claim that survived checking

An Indian approval-list entry for asenapine was located, which makes this one of the few agents here whose Indian sentence is affirmative. The entry reads: For acute treatment of schizophrenia in adults only. Asenapine Maleate Sublingual Tablets 5mg/10mg. It sits at serial number 38 of a published CDSCO approval list, dated **7 April 2011**. Note that it names two strengths where the American sublingual product carries three.

Be exact about what that date is, because the field was corrected for saying otherwise. It is the figure printed in a column headed "Date of Issue", not "Date of Approval". The field previously claimed to quote a date of CDSCO approval as printed in the source, which asserts a label the document does not carry; it now names what is actually printed, while recording that the quantity is almost certainly the approval date, since these documents are titled lists of approved new drugs. A date is labelled with the name of the column it came out of.

This is also the only Indian row in the build whose corroboration claim survived checking. Two genuinely distinct CDSCO documents carry the entry with the identical date, verified by content hash; two other agents made the same claim and it did not survive, theirs being two snapshots of one list. Establishing it needed a rendered page rather than a text extraction. One of the two documents draws no vector rules at all, so every rule-based reader refuses it; it was read from the page rendered at 200 dots per inch, showing serial 38, its acute schizophrenia indication and the

date 07.04.11, and serial 39 dated 16.04.11 below. A neighbouring row with a different date is what tells you the reader did not drift a line.

One further control sits in those lists. Entries are tagged as an additional strength when they are strength extensions rather than new approvals, and the tag is used consistently: serials 35, 36, 37 and 41 around asenapine carry it, the asenapine entry does not. That is what makes this a located approval entry rather than a permission for another dose of something approved earlier. The same tag, read carelessly on another agent in this chapter, turned a confident Indian approval date into a retraction.

IN INDIA

Write it in this order, then stop. A CDSCO published approval-list entry for asenapine maleate sublingual tablets, 5 mg and 10 mg, was located; the indication it names is "For acute treatment of schizophrenia in adults only"; the date beside it is 7 April 2011 in a column headed "Date of Issue"; it is corroborated in two content-hash-distinct CDSCO documents; and the status was true as of 1 August 2026. Do not carry the American bipolar indication or the European manic-episode indication into that sentence. Note also the direction in which this instrument fails: no India row in this build asserts a categorical negative about a drug's Indian approval status, because the published approval-list corpus is demonstrably incomplete; each asserts instead that no CDSCO approval record was located, and names what was searched. Here the corpus did display the drug, so this row needs no such hedge.

14.6 The four records side by side

Four records, one molecule, four answers, each read on a stated date.

RECORD AND APPLICATION	PRODUCT, AS THE RECORD CARRIES IT	INDICATION, IN THE RECORD'S OWN WORDS	DATE, AND WHAT KIND OF DATE
United States, NDA022117	SAPHRIS, sublingual tablet, sponsor ABBVIE; 2.5 mg, 5 mg and 10 mg as base	Schizophrenia in adults; bipolar I disorder, acute monotherapy in adults and patients 10 to 17, adjunctive to lithium or valproate in adults, maintenance monotherapy in adults	13 August 2009, approval status on the original submission
United States, NDA212268	SECUADO, transdermal system, sponsor HISAMITSU; 3.8, 5.7 and 7.6 mg per 24 hours	The treatment of adults with schizophrenia	11 October 2019, approval status on the original submission
European Union, EMEA/H/C/001177	Sycrest, holder Organon N.V., status Authorised; no dosage form is carried by this row	The treatment of moderate to severe manic episodes associated with bipolar I disorder in adults	1 September 2010, initial European Union marketing authorisation by European Commission decision
India, CDSCO approval list, serial 38	Asenapine maleate sublingual tablets, 5 mg and 10 mg; no sponsor is carried by the entry	For acute treatment of schizophrenia in adults only	7 April 2011, printed in a column headed "Date of Issue"

680 of 680	3 of 5	2 of 2	9 of 9
CHARACTERS, US ROW AGAINST LABEL	GENERIC APPLICATIONS IN THE REJECTED SET	CDSCO DOCUMENTS CARRYING THE ENTRY	FALSIFICATION CELLS PASSED

Two things about that table are as examinable as its contents. Every one of the nine falsification cells for this agent, three text instruments across three jurisdictions, returned a pass, which is why the row is authorable at all. And the European and Indian rows each record that they were built only from their own regulator's sources, with no value from another jurisdiction consulted. Isolation is built into a record, not added afterwards by an author's promise.

TRAP **Merging the two American applications into one sentence about the drug.** It reads as the tidier answer and is a false one: it produces an agent licensed for bipolar I disorder by patch, which no record here supports. The blend has already happened once in this build's own ledger. If a stem gives you a route, answer within that route's application.

14.7 What a complete answer contains

Reduced to five sentences, the item is this.

- Two American applications, two products: a sublingual tablet carrying schizophrenia and a three-part bipolar I indication, and a transdermal system carrying adult schizophrenia alone.
- Route is not packaging; it is what the indication is fastened to, and moving a patient between routes can move them outside a licence.

- Europe authorises one product for moderate to severe manic episodes of bipolar I disorder in adults, a severity qualifier the American text does not carry.
- The located Indian entry is acute schizophrenia in adults only, dated from a column headed "Date of Issue", corroborated in two genuinely distinct documents.
- Every one of those sentences carries its jurisdiction, its indication in the record's own words, its source and the date it was read, and none of them travels.

GOOD TO REMEMBER

If the patient in front of you has bipolar I disorder and the asenapine available is the product the Indian approval entry describes, the indication being treated is not the indication that entry names. That is a decision to make and document deliberately, not one made by assuming that a licence seen in an American label follows the molecule across a border.

MNEMONIC

Application, Route, Regulator, Column

Application, because two exist here and for licensing purposes they are two drugs. Route, because it is what the indication is fastened to. Regulator, because every indication sentence belongs to one and travels nowhere. Column, because a date means what its column header says it means.

Two applications, three regulators, four indications, and not one of them may be moved to where it was not issued.

15 · Lurasidone: pharmacology, metabolic profile and indications

Lurasidone is the agent on which this chapter's central hazard lands twice, in opposite directions. Clinically it is chosen for what it does not do, and the small metabolic signal in its own label, not any novelty of mechanism, is why it is here. Evidentially it defined this chapter's grounding work: its American row published a generic manufacturer's approval date under a field name meaning first approval and was corrected only on 3 August 2026, while its Indian row is a declared gap that does not recover. Where lurasidone sits in the bipolar phase algorithm belongs to the Mood disorders: pharmacotherapy notes; what follows is the whole-agent profile.

15.1 The receptor profile, and the two absences that do the clinical work

The European Summary of Product Characteristics calls lurasidone a selective blocking agent of dopamine and monoamine effects, ATC code N05AE05, and prints its affinities: 0.994 nanomolar at the dopamine D2 receptor, 0.47 at 5-HT2A and 0.495 at 5-HT7, with partial agonism at 5-HT1A at 6.38 and blockade of the alpha-2C and alpha-2A adrenergic receptors at 10.8 and 40.7. The examinable line is what is missing: the same section records that **lurasidone does not bind to histaminergic or muscarinic receptors**.

The absent histamine binding is the usual explanation for the small weight signal and the absent muscarinic binding for the lack of anticholinergic burden, but that is an inference and should be

labelled as one: the document asserts no causal link, and it states no mechanism at all for the antidepressant effect commonly attributed to the **5-HT7 antagonism** and 5-HT1A partial agonism.

Two pharmacokinetic facts change prescribing. Lurasidone is **metabolised mainly via CYP3A4** and has an active metabolite, ID-14283, itself a CYP3A4 substrate, so parent and metabolite move together whenever the enzyme is pushed. Half-life is **20 to 40 hours** with steady state within 7 days, which makes a week the honest interval before judging a dose change. The European document also records 1.5-fold higher exposure in Asian than in Caucasian subjects in a population analysis, and its posology sets no adjustment for that: worth knowing here, and not a licensed instruction to dose lower.

15.2 The metabolic profile, and exactly what the comparison licenses

In the pooled placebo-controlled adult schizophrenia studies in the American label, mean weight change was **plus 0.43 kg on lurasidone against minus 0.02 kg on placebo**, and the proportion gaining 7 percent or more of body weight was 4.8 percent against 3.3 percent. Fasting glucose moved by no more than 2.6 milligrams per decilitre in any dose arm, and total cholesterol and triglycerides fell in every arm, placebo included.

The same section carries the figures that get misused. It reports, from two of the five studies, a weight change of plus 4.15 kg for olanzapine and plus 2.09 kg for quetiapine extended-release. The clinical studies section says why the arms were there: an **active-control arm was included to assess assay sensitivity**, to show the trial could detect a drug effect at all.

0.43	4.8	22.0	4.6
MEAN WEIGHT GAIN, KG, POOLED ADULT SCHIZOPHRENIA	PERCENT GAINING 7 PERCENT OR MORE	PERCENT WITH AKATHISIA AT 120 MG	MADRS POINTS OVER PLACEBO, MONOTHERAPY

TRAP **Quoting the olanzapine and quetiapine weight figures as a head-to-head metabolic result.** They come from arms included to assess assay sensitivity, in two of five studies, with no comparative hypothesis behind them. The defensible sentence is that lurasidone's own metabolic changes were small and that two active comparators in the same programme changed weight more, stated as what the label reports and not as a ranking it establishes. Note also its statement on duration: the effectiveness of the drug for more than 6 weeks has not been established in controlled studies, so a six-week metabolic comparison is not a maintenance claim either.

15.3 What the profile costs: akathisia, movement effects and prolactin

The American label publishes the price of that profile dose by dose. Akathisia and extrapyramidal symptoms were dose-related in the schizophrenia studies: akathisia in **5.6 percent at 20 mg, 10.7 at 40 mg, 12.3 at 80 mg and 22.0 at 120 mg** against 3.0 percent on placebo, with extrapyramidal symptoms rising over the same range from 5.6 to 22.0 percent. At 160 mg akathisia was reported by 7.4 percent, 9 of 121 patients, and the label states the dose relation only up to 120 mg. That cell rests on far fewer patients and licenses no claim that akathisia falls again at the top of the range.

Prolactin is the effect most often forgotten in an agent whose selling point is metabolic. Both documents record that **lurasidone elevates prolactin through dopamine D2 antagonism**, and the European text asks that patients be counselled on gynaecomastia, galactorrhoea, amenorrhoea and erectile dysfunction. In a young adult on long-term treatment it is the conversation that decides adherence.

PITFALL

Translating "metabolically favourable" into "well tolerated". Those are different axes and the same label carries both. The ward version is a patient started at a dose chosen for speed rather than tolerability, on an agent whose **akathisia rate** runs from 5.6 percent at the lowest arm to 22.0 percent at 120 mg, and then recorded as non-adherent when they stop.

15.4 Absorption is the prescription: food, CYP3A4 and two sets of numbers

In the European food-effect study mean Cmax rose approximately 2 to 3 times and AUC approximately 1.5 to 2 times with food compared with fasting; the American label puts the same phenomenon as an AUC increase of about 2-fold and a Cmax increase of about 3-fold, and attaches a threshold the European document omits: **take with food, at least 350 calories**. A patient taking the tablet on an empty stomach is not taking a low dose of lurasidone. They are taking an unpredictable one.

WHAT THE DOCUMENT SETS	UNITED STATES PRESCRIBING INFORMATION, LATUDA	EUROPEAN SUMMARY OF PRODUCT CHARACTERISTICS, LATUDA
Unit the doses are written in	lurasidone hydrochloride: 20, 40, 60, 80 and 120 mg tablets	lurasidone base: tablets of 18.5, 37 and 74 mg, doses stated up to 148 mg, conversion printed in the document
Food	with food, at least 350 calories	once daily together with a meal, no calorie figure
Schizophrenia in adults	start 40 mg, effective 40 to 160 mg, maximum 160 mg	start 37 mg, effective 37 to 148 mg, maximum 148 mg
Strong CYP3A4 inhibitor or inducer	contraindicated; ketoconazole and rifampin named	contraindicated; ketoconazole and rifampicin named
Moderate CYP3A4 inhibitor added	halve the dose; start 20 mg, maximum 80 mg	start 18.5 mg, maximum 74 mg

The European document states in its own words that lurasidone doses of 18.5, 37, 55.5 and 74 mg are equivalent to 20, 40, 60 and 80 of lurasidone hydrochloride, so the two columns are not a disagreement: one regulator writes the salt and the other the base. A candidate who reports a maximum of 148 mg in an American stem, or 160 mg in a European one, has carried a number across a border, which is this chapter's hazard in miniature. The magnitudes: ketoconazole raised lurasidone exposure 9-fold and the metabolite 6-fold; rifampicin lowered exposure 6-fold. Grapefruit juice is avoided in both documents.

GOOD TO REMEMBER

Carbamazepine, phenytoin and rifampicin are all on the contraindicated list, all three drugs a psychiatric patient may already be taking for a reason nobody mentions at referral. This is not a caution to manage with monitoring; both regulators write it as a contraindication. Read the whole chart, not the psychotropic column, and ask directly about anti-tubercular treatment.

15.5 Bipolar depression: the size of the effect and the edges of the claim

The American label establishes this indication on three placebo-controlled six-week studies, and the effect sizes are modest. As **monotherapy** in adults, 485 patients randomised to one of two flexible ranges, the placebo-subtracted change in Montgomery-Asberg Depression Rating Scale score at week six was **minus 4.6 points, confidence interval minus 6.9 to minus 2.3, for both dose ranges**, and the label states that the higher range of 80 to 120 mg did not provide additional efficacy on average compared with 20 to 60 mg. As **adjunctive therapy with lithium or valproate** in 340 patients who remained symptomatic, the difference was minus 3.6 points, minus 6.0 to minus 1.1.

Two boundaries are examined more often than the effect size. In the label's own words, the efficacy of the drug in the treatment of mania associated with bipolar disorder has not been established: this is a bipolar-depression agent, not a mood stabiliser across phases. And fewer than 1 percent of patients in both the drug and the placebo groups developed a manic or hypomanic episode, the figure to have ready when a stem asks about switch risk. Which patient should be offered this agent rather than another is the Mood disorders: pharmacotherapy notes' question, not this one.

The paediatric headline result is a placebo-subtracted 5.7 points on the Children's Depression Rating Scale, Revised, but the two regulators print different amounts of the same study: the European text records that the primary and key secondary endpoints were not met below 15 years of age, a placebo-adjusted minus 1.8 in 10 to 14 year olds against minus 8.6 in 15 to 17 year olds, and no such age split appears anywhere in the cached American label. The regulator that authorised the paediatric indication printed less of the age detail than the one that did not, so an answer that cites only the pooled figure has taken the more permissive reading without saying so.

RULE **For lurasidone the dose that helps and the dose that harms move in opposite directions, and the label says so in two places.** The higher bipolar-depression range gave no additional efficacy on average over the lower one, while akathisia and extrapyramidal symptoms rose with dose across the schizophrenia range. Escalating because a patient has not yet responded buys a documented rise in movement adverse effects against an effect the manufacturer's own study did not find. Wait the week steady state takes before concluding a dose is too low.

15.6 Three regulators, three different sentences, one that cannot be written

The most examinable fact about this agent is that its indication is not the same in two jurisdictions and is not established from any record located in a third. The American label indicates it for schizophrenia in adult and adolescent patients aged 13 to 17, for bipolar depression as monotherapy in adult and paediatric patients aged 10 to 17, and for bipolar

depression as adjunctive treatment with lithium or valproate in adults. Section 4.1 of the European Summary of Product Characteristics authorises exactly one indication, the treatment of schizophrenia in adults and adolescent aged 13 years and over. That count is measured: the multiplicity instrument found one indication in the European primary document and one carried on the row, and the truncation instrument compared the row's 103 characters against the document's 103.

JURISDICTION	WHAT THE RECORD ESTABLISHES	AUTHORITY READ	STATUS TRUE AS OF
United States	First approval 28 October 2010: application NDA200603, brand LATUDA, original submission 1, status AP, sponsor Sunovion. The only new drug application among 21 applications for this ingredient; the other 20 are abbreviated applications	Drugs@FDA and the label archive via openFDA; the retrieved label's own application-number field reads exactly NDA200603	3 August 2026
European Union	Authorised. Initial marketing authorisation by European Commission decision on 21 March 2014, product number EMEA/H/C/002713. The register's latest Commission decision here is dated 6 May 2026 and the record does not say what it decided	Annex I section 4.1 of the Summary of Product Characteristics for the indication; the EMA register and the EPAR page for identity and status	1 August 2026, indication re-read 2 August 2026
India	No CDSCO record of a first Indian approval of lurasidone was located. What was located is a strength extension. No approval date is printable in this row	CDSCO published new-drug approval lists, 66 distinct documents, every one with an extractable text layer	2 August 2026

The European row was itself repaired on 2 August 2026: the ledger had carried the indication as adults aged 18 years and over, taken from the EMA register while citing the EPAR page as an independent check, and the register is generated from those same pages, so the corroboration was circular. And the 2026 date is not an authorisation date: that column carries the latest Commission decision, proven on this very medicine, where it reads 2026 against a 2014 authorisation. The row is also bounded by procedure: it establishes a centralised European authorisation for the named product and says nothing about products of the same substance authorised nationally, which the register does not cover by construction.

15.7 The row that was wrong, and how a generic's date wore a first approval's name

Until 3 August 2026 the American lurasidone row was false in four ways at once, each a mistake anyone can make with the same public database. Its citation listed ten abbreviated applications and no new drug application at all. Its date was 3 January 2019, a generic manufacturer's approval date, held under a field name meaning first approval, against a true first approval of 28 October 2010, an error of more than eight years. Its indication text was the label of a generic lurasidone hydrochloride tablet and not of LATUDA. And that text had been cut at 900 characters by a fixed-width slice in the grounding script.

The application endpoint paginates at whatever limit is passed, which cut the result set to ten abbreviated applications with the innovator's application absent, and the label endpoint returns a single label for a generic-name query with no guarantee that it is the innovator's. Nothing in the returned text announced the substitution, because **a generic's indication section is a near-copy of the innovator's**. Reading the indication text could never have caught it. What did was checking, for every retrieved label, whether the application number the record itself carries matched the application cited. The repaired row now passes all three falsification instruments, the truncation check reporting 974 characters on the row against 974 in the primary document; the Indian row is exempt from all three, because a declared gap carries no indication text and is a claim about a search rather than about a product.

15.8 India: a declared gap, what it rests on, and what would close it

The Indian position is a gap that is a finding rather than an absence of work. Serial number 65 of the list of drugs approved from the SND Division from 01.01.2020 to 31.12.2020 reads, in its drug-name column, "Lurasidone Hydrochloride Tablets 20mg" followed by the qualifier " (Additional Strength)", with the indication "For treatment of patients with Schizophrenia" and a date of approval of 24-12-2020. The qualifier is the entire finding. An additional-strength permission records the extension of a permission that already exists, and therefore entails a prior CDSCO approval of lurasidone at some other strength, while no such earlier entry exists anywhere in the 66-document corpus. The only date this corpus offers records a strength extension, and it must never appear in a table of approval dates. Nor can the silence be read the other way, as evidence that the extension was itself the first permission, because the source document proves an earlier one exists outside the corpus. The row was demoted from grounded on 2 August 2026 and does not recover on further retrieval of this corpus.

TRAP **Writing a categorical negative about Indian approval.** No instrument used here supports a categorical negative, and the search behind the gap is measurably incomplete: independent sensitivity in the 2024 to 2026 window is zero of one, and CDSCO's own committee minutes of 16 April 2025 record a lumateperone approval dated 26 December 2024 that appears in none of the 66 lists, although the list declaring coverage of that whole calendar year carries an entry dated 30 December 2024. Absence from this corpus is evidence, and never proof. The correct sentence names the search: no CDSCO record of a first Indian approval was located across 66 published approval lists.

IN INDIA

What would close the gap is named on the row and nothing else will do: an affirmative CDSCO or DCGI permission, a current CDSCO-approved label, or an official approval-list entry recording a first approval at any strength. Subject expert committee minutes, clinical-trial no-objection certificates, CTRI registration, brand listings, pharmacy availability and manufacturer copy are none of them marketing authorisation. That last list catches people: a resident meets lurasidone as a strip in a pharmacy long before meeting it in any gazette, and availability feels like proof. Write the search rather than the impression, and where a unit already prescribes the drug, the document it should produce is the approved label it prescribes against.

Reduced to four sentences, the agent is answerable in whatever form the question takes.

- It is chosen for the metabolic profile and paid for in akathisia, and the two are dose-related in opposite directions.
- Food is not optional: one label sets a calorie threshold, the other does not, and both make the dose meaningless without a meal.
- CYP3A4 is the interaction axis, and strong inhibitors and inducers are contraindications rather than cautions.
- The indication set is American and broad, European and narrow, and Indian and not established from any record located.

MNEMONIC

Food, Enzyme, Dose, Salt

The four things that go wrong with lurasidone: given without a meal, given with a strong CYP3A4 inhibitor or inducer, escalated past the dose its own evidence supports, and its milligrams quoted from the wrong regulator's arithmetic.

Lurasidone's dose, its indication and its approval date each mean something different in each of the three jurisdictions, and the Indian date is the one no document read here contains.

16 · Brexpiprazole: pharmacology and indications

This is the agent on which a candidate is likeliest to write a true sentence about the wrong country. Brexpiprazole is a dopamine partial agonist whose licensed scope differs at every border read for this chapter: three indications in one jurisdiction, one in each of the other two. The marks separate on pharmacology, on jurisdiction, and on the comparison with cariprazine. Its position in the depression-augmentation ladder, beside lithium, thyroid hormone and the other atypical agents, belongs to the Mood disorders: pharmacotherapy notes.

16.1 Partial agonism is a claim about intrinsic activity

The word carrying this agent is **partial agonist**, and it states **intrinsic activity** rather than affinity. A partial agonist binds the receptor and produces a submaximal response, so what it does

depends on the tone it meets. Where dopamine transmission is high it competes with the endogenous full agonist and lowers signalling; where transmission is low it supplies a floor of activity instead of removing the little that was there. That asymmetry is the rationale of the class: damp the mesolimbic excess without emptying the nigrostriatal and tuberoinfundibular pathways to the same degree.

Two consequences follow. The burdens of full D2 blockade, extrapyramidal signs and prolactin elevation, are expected to be lighter rather than absent, because partial agonism reduces the depth of blockade and does not abolish it. And the class carries a signature adverse effect of its own rather than the antagonists': activation, inner restlessness and **akathisia**, which is residual agonist drive at D2 seen at the bedside. A candidate who writes that partial agonists are simply better tolerated has stated the first and lost the second.

16.2 Where brexpiprazole sits between aripiprazole and cariprazine

Hold brexpiprazole as a deliberate re-engineering of aripiprazole rather than a new idea: a partial agonist at D2 and at 5-HT1A and an antagonist at 5-HT2A, carrying lower intrinsic activity at D2 than aripiprazole and weighted more heavily toward its serotonergic and alpha-adrenoceptor arms. The proprietary class phrase attached to it, serotonin-dopamine activity modulator, is manufacturer language and not a regulatory category; describe the profile and leave the phrase alone.

The re-engineering exists for a clinical trade, and the trade is the answer: less of the activation and akathisia that limit aripiprazole, paid for with more sedation and more weight gain, because the serotonergic and adrenergic arms do more of the work. Hold that as a direction and not a magnitude: no head-to-head figure is printed here, because no source read for this fragment carries one, and a remembered effect size is worth less than a correctly stated direction.

Cariprazine separates from both on two features: it is D3-preferring at clinically relevant occupancy rather than merely D3-active, and it carries a long-acting active metabolite, so its offset after a dose change is far slower than the parent compound suggests. An adverse effect does not resolve on the timescale a resident expects, and a titration decision expresses itself over weeks. Brexpiprazole carries no comparable metabolite story, which is why examiners ask about the two together.

RULE **The mechanism is shared, the profile is not, and the licence is neither.** Three layers travel with a partial agonist: what it does at the receptor, what that costs the patient, and what a named regulator has permitted in a named jurisdiction on a named date. Each layer is sourced in its own right, and a sentence that slides from one to the next has stopped being checkable.

16.3 The United States licence, read whole, including the limitation

The American label is the widest of the three. It carries adjunctive treatment to antidepressants for major depressive disorder in adults; treatment of schizophrenia in adults and pediatric patients ages 13 years and older; and treatment of agitation associated with dementia due to Alzheimer's disease. Three indications, three populations, one molecule.

Then the clause a truncated reading loses. The label attaches a **Limitation of Use**: REXULTI is not indicated as an as needed, "prn", treatment for agitation associated with dementia due to Alzheimer's disease. Read what that limits. It does not withdraw the dementia-agitation indication; it forbids the way a busy ward is most tempted to use it, as a rescue dose written up in advance. The licensed use is scheduled treatment of a syndrome, not a chemical response to an incident.

The status was read from the openFDA structured product label archive, brand REXULTI, label effective 6 May 2026, whose application-number field returns NDA205422, the only new drug application among 17 applications for this active ingredient, the other 16 being abbreviated applications. First approval is recorded as NDA205422, sponsor Otsuka, original submission 1 carrying an approved status dated **10 July 2015**, and the reading was status true as of 3 August 2026. Jurisdiction, indication in the source's words, named application, dated reading: four fields, and an answer missing any of them is an assertion rather than a record.

TRAP **Quoting the three American indications as though they were the drug's indications.** They belong to one label in one country; the other two jurisdictions read here each carry exactly one. The failure is not ignorance of the other licences; it is the silent generalisation, writing "brexpiprazole is indicated for agitation in Alzheimer's disease" with no country attached. The label text is the current wording as filed with one regulator, and it is not a statement about any other jurisdiction and must never be carried across one. Attach the country to every indication sentence and the trap cannot close.

16.4 Europe: the same molecule, one indication

The European licence is narrower, and short enough to memorise. Section 4.1 of the Summary of Product Characteristics reads that RXULTI is indicated for the treatment of schizophrenia in adults and adolescents aged 13 years and older, and the falsification pass counted one indication in that section against one on the row. No adjunctive-depression indication appears in it. The sentence to write is that the European authorisation read here covers schizophrenia and names nothing else, which is a statement about a document rather than a claim about what exists elsewhere.

The identity fields matter as much as the text. The product is Rxulti, European Medicines Agency product number EMEA/H/C/003841, with Otsuka Pharmaceutical Netherlands B.V. as **marketing authorisation holder**, medicine status Authorised; the initial European marketing authorisation, by European Commission decision, is dated **26 July 2018**, and the most recent Commission decision on it is dated 9 July 2026. The reading was true as of 1 August 2026.

Two disciplines sit behind that paragraph. The first is which document is the authority: the indication text is taken from Annex I section 4.1 of the product information, because the agency's published medicines register is a machine dump generated from the same assessment-report pages and cannot corroborate them independently. Two views of one source are one source. The second is scope. An entry in that register establishes authorisation for the named product through the **centralised procedure**, and says nothing about other products containing the same

active substance that may be authorised nationally through decentralised or mutual-recognition routes, which the register does not cover by construction.

16.5 India: the row whose column header genuinely says approval

The Indian entry is the recent one, and it is the one an Indian examination cares about. The Central Drugs Standard Control Organisation entry read here states that brexpiprazole is indicated for treatment of schizophrenia, as brexpiprazole tablets **0.25, 0.5, 1, 2, 3 and 4 mg**. Those are permitted strengths and not a dose recommendation; no prescribing dose for this agent is printed here, because no source read for this fragment carries one, and the augmentation dose belongs with the ladder in the Mood disorders: pharmacotherapy notes.

The date and its label are the transferable part. The entry is dated **16 February 2026**, printed in a column headed **Date of Approval**, which makes this the only grounded Indian row in the chapter whose source header genuinely says approval rather than issue. Elsewhere the honest field name is the printed one: a date under a column headed date of issue must not be relabelled a date of approval by whoever quotes it. The provenance is entry serial number 1 in the list of new drugs approved in the year 2026 to date, within a corpus of 66 distinct published approval documents, with a separate bulk-drug permission for the substance dated 24 February 2026.

How strong is that row? Stronger than most. Its corroboration is one running list captured in two snapshots and not two distinct documents, a description corrected on 2 August 2026 from the stronger claim; being about five months old at retrieval it is the most recency-sensitive row in the Indian column, so it was re-fetched live and the source list came back byte-identical, which is the best evidence this row has. Two snapshots of one list is weaker corroboration than two independent documents, and saying so is what makes the row trustworthy rather than what weakens it.

TRAP **Turning a located approval into a categorical negative.** The Indian entry names schizophrenia, and the temptation is to write that brexpiprazole is not approved in India for anything else. Do not write that; nothing here supports it. The published approval-list corpus is demonstrably incomplete, so an absence in it is evidence and never proof, and the defensible sentence names what was located and what was searched. The mirror error runs the other way: a brand listing, a pharmacy shelf or a manufacturer's page is not a regulatory approval and must never be quoted as one.

IN INDIA

Write it in this order. A CDSCO new-drug approval record for brexpiprazole was located, dated 16 February 2026, under a column headed Date of Approval, naming schizophrenia and six tablet strengths, with a separate bulk-drug permission a week later. That entry establishes what it says and nothing wider, and it is not an availability or a price claim, neither of which any source read here establishes. For a patient or a family: an approval on a national list is not the same as a drug on the nearest shelf at a price the household can meet, and that has to be checked locally.

16.6 Adjunctive depression, and the comparison with cariprazine

This is where the agent earns its place, and where the cariprazine comparison stops being trivia. In the United States both drugs carry it: brexpiprazole's label reads adjunctive treatment to antidepressants for major depressive disorder in adults, and cariprazine's reads **adjunctive therapy to antidepressants** for the treatment of major depressive disorder in adults. Two dopamine partial agonists, one indication, so the choice between them is made on profile rather than on licence.

Cross the Atlantic and the symmetry breaks. The European product information for RXULTI names schizophrenia in adults and adolescents aged 13 years and older, and cariprazine's names the treatment of schizophrenia in adult patients. Neither European licence read here carries the adjunctive-depression indication, and brexpiprazole's is the one extending below adult age. An answer that reproduces the American adjunctive indication for a European or Indian patient has made the commonest error on this topic.

3	1	1 of 17	9 of 9
INDICATIONS, UNITED STATES LABEL	INDICATIONS, EUROPEAN SECTION 4.1	US APPLICATIONS THAT ARE AN NDA	FALSIFICATION CELLS PASSED

JURISDICTION	INDICATION, IN THE SOURCE'S WORDS	SOURCE AND DATE
United States	Adjunctive treatment to antidepressants for major depressive disorder in adults; schizophrenia in adults and pediatric patients ages 13 years and older; agitation associated with dementia due to Alzheimer's disease, plus a Limitation of Use against as-needed dosing for the dementia indication	Structured product label, application NDA205422, effective 6 May 2026; true as of 3 August 2026
European Union	The treatment of schizophrenia in adults and adolescents aged 13 years and older	Product information Annex I section 4.1, EMEA/H/C/003841; true as of 1 August 2026
India	Treatment of schizophrenia, tablets 0.25, 0.5, 1, 2, 3 and 4 mg	CDSCO published approval list, entry 1, 2026 list; true as of 1 August 2026

16.7 How a status sentence is checked, and the defect on this very row

This row was tested rather than trusted. Three text instruments were run against the primary document for each of the three jurisdictions and all nine cells passed: that the row text is verbatim from and attributed to the primary document; that it is not the primary text cut short; and that the count of indications carried matches the count in the document. The American cell measured 985 characters against 985 and five indications against five, the European 105 against 105 and one against one, and the Indian cell was carried in full. The Indian cell was located

geometrically, at entry 1 on page 1 within a named indication column, with two independent readers agreeing with the drawn table rules.

Now the defect this row carried, because its shape is the lesson. Until 3 August 2026 the American entry cited an abbreviated application, a generic's application, ahead of the innovator NDA205422, and its indication text had been truncated at 900 characters in mid-word by the script that built it. The date was never wrong. Sit with that. A right number under a wrong citation is more dangerous than a wrong number, because nothing about it looks incorrect and every downstream reader inherits a claim its stated evidence does not support. The check that catches it is not a check on the number: it asks whether the cited document could have established the fact, and an **abbreviated application** is a copy of an innovator's licence, never the origin of a first approval.

PITFALL

Completing a sentence that stops in the wrong place. The truncation above cut a label mid-word; its sibling defect cuts text so that it ends in a clean full stop the source never had, and nothing on the page then signals the loss. Both are invisible to a reader willing to finish the sentence from memory. When a quoted indication reads as though a clause is missing, say so and go back to the primary document. Every repair in this chapter's American column began with someone refusing to fill a gap and saying it aloud.

16.8 What a complete answer contains

Reduced to five moves, the item is answerable in any form.

- State partial agonism as intrinsic activity, then what it buys, lighter blockade burden, and what it costs, activation and akathisia.
- Separate the three partial agonists by discriminator: intrinsic activity at D2, the serotonergic weighting, and the long active metabolite that is cariprazine's and not brexpiprazole's.
- Attach a country, a source and a date to every indication sentence, and never carry one across a border in either direction.
- Quote the American limitation on as-needed use for dementia agitation, the clause most often lost and the one that changes ward practice.
- State an Indian record as what was located and where, name the printed column header, and refuse the categorical negative no instrument here supports.

GOOD TO REMEMBER

The commonest real-world error here is not choosing the wrong partial agonist. It is writing a rescue prescription for agitation in dementia against a licence that specifically excludes as-needed use, and then defending it with an indication that does exist. The indication and the manner of use are two separate permissions, and the second is where this label draws its line.

MNEMONIC

Tone, Trade, Territory

Tone: partial agonism acts on what the receptor is already seeing. Trade: less activation, more sedation and weight. Territory: three indications in one country, one in each of the other two, each with its own source and date. An indication sentence with no territory in it is not yet an answer.

Partial agonism is intrinsic activity, the trade is akathisia against weight, and no indication is true until a country, a source and a date are attached to it.

17 · Cariprazine: pharmacology and indications

One molecule, three national records, and four, one and two indications. The pharmacological question is what a D3-preferring partial agonist does that a full antagonist cannot, and what a long-lived active metabolite does to every decision after a dose change. The regulatory question is why it carries four licensed indications in one jurisdiction, one in the second and two in the third, all three grounded and each read on its own day. Marks are lost in three places: candidates upgrade a receptor rationale into a licensed claim, they reason about dose changes on a timescale this drug does not obey, and they carry the bipolar-depression indication across a border where no record read here places it. Phase matching, mixed states and where this agent sits in a bipolar algorithm belong to the Mood disorders: pharmacotherapy notes and are not rebuilt here.

17.1 Partial agonism, and the D3 preference that is the design claim

Begin with what the regulator wrote. The current United States label calls the product an atypical antipsychotic and names no receptor and no subtype. Every receptor sentence below is therefore class pharmacology taught as rationale, carried by no claim row in this build, and must never be written as though a regulator endorsed it.

A **partial agonist** occupies the receptor but produces a submaximal response even at full occupancy, and its intrinsic activity decides the net effect: at high dopamine tone it replaces a stronger signal with a weaker one and behaves as an antagonist, at low tone it supplies signal and behaves as an agonist. That is why marked prolactin elevation is not the expected consequence of adequate occupancy, and why the characteristic adverse effects are activating rather than sedating. It also explains a prescribing hazard: adding a partial agonist to a full antagonist is not additive blockade, because the newcomer competes for occupied receptors and substitutes its weaker signal, so a stable patient can destabilise during cross-titration.

The D3 claim distinguishes this agent from the other dopamine partial agonists here, and is most often overstated. Cariprazine is **D3-preferring**: in laboratory binding work its affinity for D3 exceeds that for D2, and D3 receptors sit in limbic territory tied to motivation and reward, hence the rationale in negative and anhedonic presentations. Two cautions belong beside it. Affinity in vitro need not become preferential occupancy in a living brain, where endogenous dopamine itself binds D3 far more potently than D2. And a rationale is not an outcome: no trial-level

evidence lane exists for this agent in this build, so no negative-symptom result, effect size or comparator is printed here.

RULE **A receptor preference is a hypothesis about mechanism; an indication is a sentence a regulator wrote.** They carry entirely different warrants. Write the pharmacology as rationale, write the indication in the regulator's own words with jurisdiction, source and date attached, and never let the first widen the second.

17.2 The long metabolite, and why this drug's clock runs in weeks

Cariprazine is metabolised to two **active metabolite** species, and the didesmethyl form is eliminated far more slowly than the parent. The relevant quantity is therefore the total active moiety, which accumulates towards steady state over weeks, not the two or three days a clinician's reflex assumes. No half-life figure is printed, because no claim row in this build carries one and a number written from memory is the defect this chapter exists to teach against. The shape of the curve is the examinable part.

Four consequences follow, each a clinical decision rather than a pharmacokinetic fact. A dose increase judged at one week is judged too early, its plasma consequence still rising. A reduction made for an adverse effect will not relieve it promptly, so a clinician seeing no change within days is tempted into a second reduction that compounds the first. Discontinuation does not produce a drug-free patient on the timescale the chart implies, which matters for a planned switch, for pregnancy planning, and for a new adverse effect a fortnight after stopping. And washout arithmetic learned on a short-acting agent gives the wrong answer.

GOOD TO REMEMBER

Write the review interval to the drug's clock rather than the clinic's. Where the active moiety is still accumulating, an unchanged picture at one week is not treatment failure and a new adverse effect at six weeks is not a new drug. Make one change and wait longer than feels natural: the commonest error is a rapid sequence of adjustments that leaves the eventual picture uninterpretable.

17.3 What the mechanism predicts, and the figures this build does not hold

The adverse-effect profile is predictable from partial agonism. Activating effects dominate: **akathisia** is the signature, restlessness and insomnia accompany it, both dose-related. Other extrapyramidal features occur but do not lead, and nausea and dyspepsia are common early. Because adequate occupancy is reached without complete blockade, marked prolactin elevation is not expected, and the metabolic burden is described as modest rather than absent.

Akathisia is where the two sections above meet: most likely to bring a patient back, most likely to be misread as agitation belonging to the illness, slowest to answer a dose reduction because of the long active moiety. A patient told the restlessness will settle in a day or two has been given a promise the pharmacokinetics cannot keep.

Now the boundary. Every sentence in this section is receptor-level and class-level reasoning. No adverse-event rate, no comparative metabolic figure, no weight-change number and no

comparative akathisia incidence is held in this build for this agent, because no evidence lane was grounded and the retrieved label section covers indications only.

PITFALL

Quoting a comparative metabolic or akathisia figure from memory. It sounds like something a monograph carries, which is why the fabricated version passes unnoticed. Two defensible sentences exist: the mechanism predicts an activating rather than a metabolic profile, and the size of that difference is a question for a named trial. Supplying a percentage turns a plausible expectation into a fabricated result.

17.4 The United States record, read field by field

This row carries the widest indication set, so it is the one material leaks from. The label indicates the product for the Treatment of schizophrenia in adults and pediatric patients 13 years of age and older; for Acute treatment of manic or mixed episodes associated with bipolar I disorder in adults and pediatric patients 10 years of age and older; for Treatment of depressive episodes associated with bipolar I disorder (bipolar depression) in adults; and as Adjunctive therapy to antidepressants for the treatment of major depressive disorder (MDD) in adults. Four indications, two reaching below adulthood at different ages, the third the bipolar-depression licence distinguishing this agent among the dopamine partial agonists.

The source is named and narrow: the openFDA Drugs@FDA interface, application NDA204370, brand VRAYLAR, sponsor AbbVie, with the current label effective 18 December 2025. The document is a **structured product label**, and the status true as of date on this row is 3 August 2026. The row cautions about its own instrument: openFDA is an agency-operated mirror of Drugs@FDA and the label archive rather than the Drugs@FDA application itself, and the indication wording is current label text as filed, which is not a statement about any other jurisdiction and must never be carried across one.

Then the date field, where this ledger has been burned. The earliest approval-status date in the record is **17 September 2015**, named for what it holds, not what a reader wants it to mean. It also shows a single application returned for this agent, thirteen approval-status submissions within it, and the most recent dated 18 December 2025, the label date quoted above. Whether the earliest of thirteen is a first marketing approval, the record does not say; elsewhere in this build a generic's approval date was published under a first-approval field name and retracted.

4	1 of 1	1125 of 1125	13
INDICATIONS ON THE CURRENT US LABEL	APPLICATIONS RETURNED FOR THIS AGENT	CHARACTERS, ROW AGAINST PRIMARY	APPROVAL-STATUS SUBMISSIONS IN THE RECORD

This row was tested, not trusted, and it had a real defect. Until 3 August 2026 its indication text was cut at 900 characters, mid-word, by a fixed slice in the grounding script that left no marker; the cited application was correct throughout, so the row was truncated only and never misattributed. The repaired row then passed three instruments against the primary document: the row text is verbatim, it is not the primary text cut short, at 1125 characters against 1125, and

the count of indications carried matches the count in the document at five against five. That last figure is its parse of a label printing the list twice, in the numbered section and the highlights summary; the teaching is that the counts agree, not that five things are licensed.

17.5 The European Union: one indication, and a corroboration that was circular

The European record is grounded and narrow. The authorised indication, in full: Reagila is indicated for the treatment of schizophrenia in adult patients. One indication, adults only, nothing on mania, bipolar depression or adjunctive use. The instrument: Annex I Summary of Product Characteristics, section 4.1 Therapeutic indications, read from the cached product information and re-read verbatim on 2 August 2026.

Product identity and status come from a separate place and corroborate identity only: the agency medicines register lists Reagila, product number EMEA/H/C/002770, status Authorised, and an independent recheck of the assessment-report page returned the same product number, status and authorisation date. The **marketing authorisation holder** is Gedeon Richter. The initial marketing authorisation, a European Commission decision, is dated **13 July 2017**; the most recent Commission decision on the file is dated 8 January 2026.

One correction here is worth more than the dates above. The indication was originally taken from the medicines register, with the assessment-report page cited as an independent second instrument. The register is generated from those same pages, so the corroboration was circular and blind to the very failure it was meant to catch, a stale summary layer. Section 4.1 of the annex is the genuinely independent instrument and is now the authority for indication text. Two renderings of one source are one source; citing a summary and its own summary as agreement checks nothing.

The scope limit travels with it. This establishes marketing authorisation for the named product through the European Union **centralised procedure**, and says nothing about other products containing the same active substance that may be authorised nationally through the decentralised or mutual-recognition procedures, which this register does not cover by construction.

17.6 India: the located CDSCO entry, and the two indications it carries

The Indian row is the one an Indian examination cares about, and here it is affirmative rather than absent. A CDSCO published approval list carries the entry, whose cell reads whole: For the treatment of Schizophrenia in adults. Acute treatment of manic or mixed episodes associated with bipolar I disorder in adults. Cariprazine hydrochloride bulk and Cariprazine capsules 1.5mg/3mg/4.5mg and 6mg. Two indications, both adult, and a product cell naming the **bulk drug** and the capsules at **1.5mg/3mg/4.5mg and 6mg**. Those strengths are a formulation fact from a name cell, not a dose; no dose, titration or schedule for this agent is held in this build.

The locator is exact, the date precisely labelled. The entry is serial number 15 in the 2021 CDSCO approval list, one document in the 66-document published corpus retrieved on 1 August 2026, and the date carried is **16 July 2021**. It is not called an approval date, and the reason is instructive: the field records the date printed in the column headed **Date of issue**, after an earlier version claimed to quote a label the document does not carry. It is almost certainly the approval

date, these being titled lists of approved new drugs, but the field now names what is printed. Corroboration is one CDSCO document, the date confirmed against the raw column layout rather than the flattened text line.

Now the discipline governing what the entry omits. It names schizophrenia and acute manic or mixed episodes, and neither bipolar depression nor adjunctive use in major depressive disorder. That is a statement about the record, not the country. The published approval-list corpus is demonstrably incomplete: every drug it is known to find sits in a document dated 2023 or earlier, its measured sensitivity in the 2024 to 2026 window is nil, and its positive-control tally may not warrant a gap on any agent approved after 2023. No India row in this build asserts that a drug is not approved in India; each asserts that no CDSCO approval record was located, and names what was searched. The same restraint applies to an indication absent from a located entry.

IN INDIA

Write it in this order. A CDSCO approval entry for cariprazine was located, dated 16 July 2021 under the printed heading Date of issue, covering the bulk drug and four capsule strengths, with two adult indications: schizophrenia, and acute manic or mixed episodes of bipolar I disorder. It carries no bipolar-depression indication, a statement about the retrieved record and not about Indian practice. Do not write that any use is unapproved in India: the corpus that would have to be complete for that sentence to be safe is demonstrably not.

TRAP **Carrying the bipolar-depression indication across a border.** It belongs to one record only: the American label carries it, the European annex carries schizophrenia in adults alone, the Indian entry schizophrenia and acute mania. Writing that cariprazine is indicated for bipolar depression, without naming the jurisdiction, is true in one place, unsupported in the second and beyond the located record in the third. The paediatric ages travel no better: 13 and 10 years are American label figures appearing in neither other record.

17.7 How that Indian cell was certified, and what a certified cell is

The entry sits in a scanned document, so reading a cell out of it is an extraction problem before a clinical one. Of three geometric readers in this build, two could not be used here: one refuses the merged block spanning serial numbers 14 and 15, the other is refused across the whole document because its declared page-level invariant is falsified there. The cell was certified instead by agreement between the reader working from the drawn table rules and one rendering the page as pixels and reading it by eye. Two readers of genuinely different kinds agreeing is what certification means; two text readers agreeing is one reading counted twice.

Certification earned its cost at once. The moment the cell became certifiable, the completeness instrument failed the row: the published text had dropped the bulk-drug leg of the name cell, which reads Cariprazine hydrochloride bulk and Cariprazine capsules. A comparable entry on another agent in this chapter retains its bulk leg, so the loss was an inconsistency rather than policy, and the leg is restored. Repaired, the three instruments pass against a locator giving document, page, serial number and the pixel bounds of the indication column and row band: the

text is verbatim, it is not the primary cut short, at 215 characters against 215, and the 133-character indication cell is carried in full.

The lesson is the order of events: the defect was invisible while the cell was unreadable and appeared in the hour it became readable. An unaudited row is not a clean row, and defect counts rise when the instrument starts working.

17.8 The three records side by side, and what a complete answer contains

Three regulators, three reading dates, three indication sets.

JURISDICTION	WHAT THE RECORD HOLDS	THE SENTENCE THAT MAY BE WRITTEN	THE SENTENCE THAT MAY NOT
United States	Four label indications: schizophrenia from 13 years, acute mania or mixed episodes from 10 years, bipolar depression in adults, adjunctive therapy in major depressive disorder in adults, NDA204370, label effective 18 December 2025, read 3 August 2026	Each indication whole, with application, label date and reading date	Any wider indication, and the earliest date called a first approval
European Union	One authorised indication, schizophrenia in adult patients, from section 4.1 of the annex; authorised centrally from 13 July 2017	The indication verbatim, with the centralised-procedure scope limit	Any American indication transferred, and any claim about national authorisations this register cannot see
India	A located CDSCO entry carrying two adult indications, the bulk drug and four capsule strengths, serial number 15, dated 16 July 2021 under the heading Date of issue	The entry whole, with its serial number, document and the heading its date came from	Any negative inference from what it omits, and approval substituted for the printed heading

Reduced to five sentences:

- Partial agonism explains the shape of benefit and harm: submaximal signal at high tone, supplied signal at low tone, an activating adverse-effect profile.
- The D3 preference is a laboratory affinity ranking offered as rationale for negative and anhedonic presentations, neither outcome nor licensed claim.
- The long-lived active metabolite moves every clinical timescale outward: dose changes judged late, adverse effects resolving slowly, stopping that is not clearing.
- The indication set differs by jurisdiction and the bipolar-depression licence sits in one record, so no status sentence may be written without its jurisdiction, source and date.

- What a record omits is a fact about that record, and on the Indian corpus it cannot be converted into a statement about the country.

MNEMONIC

Preference, Persistence, Place

Preference is the D3 story, a binding ranking offered as rationale and never as licence. Persistence is the long active moiety: accumulation, delayed judgment, slow offset. Place is the jurisdiction every indication sentence must name and never leave, and why four, one and two are all correct answers here.

The receptor preference is a rationale, the metabolite makes the clock run in weeks, and the indication count is four, one or two depending on which regulator you are quoting.

18 · Lumateperone: pharmacology and indications

One capsule acts on serotonin, dopamine and glutamate at the same time, and the breadth of the licence follows from that rather than from any single receptor. Lumateperone is examinable for two unrelated reasons, and preparing one loses half the marks. The pharmacology leaves the dopamine-occupancy logic every antipsychotic algorithm is built on, which is what makes one molecule plausible across schizophrenia and bipolar depression. The regulatory record is the sharpest case in this chapter: three national records of one drug in three states, the Indian one neither an approval nor an absence. Receptor principles for the serotonin and dopamine systems belong to the Psychopharmacology I: principles and core drug classes notes and are not rebuilt here.

18.1 One molecule across three transmitter systems

The mechanistic account has four legs, taught as pharmacology rather than as anything a regulator has endorsed.

- Potent antagonism at the 5-HT2A serotonin receptor, the highest-affinity action of the molecule.
- A phase-dependent action at the dopamine D2 receptor, partial agonism presynaptically with antagonism postsynaptically, damping transmission where it is excessive rather than subtracting it everywhere.
- Inhibition of the serotonin transporter, an action shared with antidepressants rather than antipsychotics.
- Indirect enhancement of glutamatergic transmission in prefrontal circuits, dependent on the dopamine D1 receptor and expressed through NMDA and AMPA currents.

A molecule doing all four has a route into psychosis and a route into depressed mood at once, which is why the licence in 18.2 stops looking like three unrelated permissions. That is a rationale, not a demonstration: no trial evidence for this agent is held in this build, so no trial name, sample size, endpoint or effect size appears here.

The licence names none of the four legs. The current United States label calls the product **an atypical antipsychotic** and stops: no receptor, no subtype, no affinity ratio, no occupancy figure. Marks go in a predictable direction, by promoting the mechanism into the licence or by attaching an occupancy percentage to a regulatory sentence. No occupancy or affinity figure is printed here, because none is grounded in this build, and a number recalled from a review and set beside a licence is a fabricated citation whatever its provenance.

RULE **A mechanism explains why a licence might be wide; it never widens one.** Write what the molecule does, then separately what a named regulator has permitted, in that regulator's words, with the date of the reading. The two sentences must stand apart: one is pharmacology that travels everywhere, the other a permission that stops at a border.

18.2 The United States: one application, three indications, quoted whole

Exactly one jurisdiction carries a grounded status row here, and its fields are what make a claim checkable. The source is named and narrow: the openFDA interface to Drugs@FDA, one application, NDA209500, with the current structured product label effective 20 June 2025, and the status was true as of 1 August 2026. The indication text follows, with the label's pointers to its own clinical-studies sections removed and nothing else changed.

CAPLYTA is indicated for: Treatment of schizophrenia in adults. Treatment of depressive episodes associated with bipolar I or II disorder (bipolar depression) in adults, as monotherapy and as adjunctive therapy with lithium or valproate. Adjunctive therapy with antidepressants for the treatment of major depressive disorder (MDD) in adults.

Three permissions, and they are not parallel.

LICENSED INDICATION, UNITED STATES	MAY BE GIVEN ALONE	MAY BE GIVEN WITH
Treatment of schizophrenia in adults	No partner drug named either way	No partner drug named either way
Depressive episodes associated with bipolar I or II disorder (bipolar depression) in adults	Yes, stated as **monotherapy**	Also stated as **adjunctive therapy** with lithium or valproate
Major depressive disorder in adults	No such wording appears	Adjunctive therapy with antidepressants only

The asymmetry is the mark. Bipolar depression is licensed both ways; major depressive disorder one way only, as an addition to an antidepressant. The bipolar row covers bipolar I and bipolar II together, which much of the literature a candidate has read does not. Note what the schizophrenia row does not say: an adult indication, no adolescent extension, no line of therapy, no comparator, no place in any algorithm.

Then the field most often misreported. The earliest approval-status date for this application is **20 December 2019**, and the field is named **the earliest approval-status date in the record**. Whether that is a first marketing approval, a supplement or something else, the record read here does not say. This ledger has already retracted one drug's date after publishing it under a field

name that promised more than the field held. Print a date with the name of the field it came from.

Two disciplines travel with the American row. The interface is an agency-operated mirror of Drugs@FDA and of the label archive rather than the Drugs@FDA application itself, and the indication wording is current label text as filed, which is not a statement about any other jurisdiction and must never be carried across one. And the row was tested rather than trusted. Three text instruments ran against the primary document, checking that the row text is verbatim, that it is not the primary text cut short and closed so that nothing signals the loss, and that the counts of indication statements match. All three passed, at 848 characters against 848, and five indication statements against five. Five rather than three counts the label's highlights repeat of the same list.

GOOD TO REMEMBER

Licensing bipolar depression both as monotherapy and as an addition to lithium or valproate answers the question the ward asks, which is whether the existing mood stabiliser has to come off. On this wording it does not. What the licence does not answer is whether adding this drug beats optimising what is there, because no comparative, sequencing or head-to-head result for this agent is held in this build. Answer the licensing question from the label; refuse the comparative one aloud.

TRAP **Carrying the monotherapy permission across from bipolar depression to major depressive disorder.** The wording gives bipolar I and bipolar II depression as monotherapy and as adjunct; it gives major depressive disorder as adjunctive therapy with antidepressants and in no other form. Writing that the drug is licensed for depression, unqualified, merges two indications with different permissions and different partner drugs, and usually exports an American permission to whichever country the examiner had in mind.

18.3 The European Union: an absence in a register that records refusals too

A declared gap is a statement about a search and not about a product, so its strength is its instrument's strength. The instrument is the European Medicines Agency's own published medicines register, generated by the agency on 1 August 2026, downloaded whole and scanned across all 39 columns of all 2730 medicine rows, where this active substance returns zero rows under any of its names.

What lifts that above a shrug is the range of outcomes the register carries. It records Authorised, Refused, Application withdrawn, Withdrawn, Lapsed, Expired, Revoked, Suspended and Opinion outcomes, so the absence means no central application of any kind was ever concluded, including one refused or abandoned partway. The missing document is an agency public assessment report for the substance or a European Commission implementing decision, and the phrase approved internationally is not a permitted substitute for a named regulator.

Now the limit, where an honest answer separates itself from a confident one. The register covers the European Union **centralised procedure** only. A product authorised nationally in one or more member states through the decentralised or mutual-recognition procedure would never

appear in it, so absence here means no centralised procedure of any kind is on record, which is narrower and weaker than not available anywhere in the Union. That limit sits on every European gap row and does not bite equally on all; on this agent it is graded as biting moderately. Two weaknesses travel with the finding. The confirming second instrument ran on all eight grounded European rows and on none of the eight gap rows, so no European gap here has independent corroboration, and the assessment-report page is in any case a second rendering of the same source, because the register is generated from those pages.

The defensible European sentence is short and worth writing whole: no European Union centralised procedure of any kind is on record for lumateperone as of 1 August 2026, which does not exclude a national authorisation this instrument cannot see.

18.4 India: a committee that recites an approval no list prints

The Indian entry is the one an Indian examination cares about, and it is the only row in this chapter that is neither a grounded approval nor a declared gap. Its state is partial, on a supportive source only. One affirmative document exists, a committee minute rather than a permission. Minutes of the CDSCO **subject expert committee** for neurology and psychiatry, dated 16 April 2025, record verbatim that the committee noted that drug Lumateperone Capsules 42mg is already approved in the country on 26.12.2024. The date recited is **26 December 2024**, and the strength named, **42 mg**, is part of the product designation in that minute and not a dosing instruction; no dosage text for this agent is held in this build.

What the row does not hold is a **marketing authorisation**. Read the next sentence twice, because collapsing it is the most expensive error on this item. Subject expert committee minutes are a supportive source only, so the row licenses the statement that CDSCO committee minutes of 16 April 2025 recite such an approval, and it does not license the statement that CDSCO approved lumateperone on 26 December 2024. Two claims about different things, one about a document and one about a regulatory act, and only the first is in evidence. The second would rest on the Drugs Controller General of India, whose permission was searched for and not located.

The primary record was sought in three forms and found in none. The approved indication wording was not located at all, so the row records the fact of an approval being recited and not its indication. No CDSCO permission and no approval-list entry appears in the 66 published approval-list documents. A run-time scan of all 121 extracted CDSCO texts found the drug's name in none of them.

The list that should have carried it demonstrably covers the date. The list declaring calendar-2024 coverage carries entries dated 17 December 2024 and 30 December 2024, and declares its own range in its running header as covering 1 January 2024 to 31 December 2024, so it is not truncated before 26 December. Two corrections to that argument sit in the row and are worth learning as technique. An earlier version cited entries dated 17, 30 and 31 December 2024, but 31 December occurs only inside the page header and is not an entry; the declared coverage range is the stronger evidence anyway, because it excludes truncation by the document's own statement rather than by inference from its last row. A second leg, that the list carries roughly 15 numbered

rows, was withdrawn as false: there is no single calendar-2024 list, the corpus holds five snapshot vintages across two CDSCO series, and the one declaring full-year coverage carries 30 recoverable rows to a highest serial of 31.

IN INDIA

Write it in this order and stop. CDSCO subject expert committee minutes of 16 April 2025 recite an approval of Lumateperone Capsules 42 mg in the country on 26 December 2024; committee minutes are a supportive source and not a marketing authorisation; the permission, an approval-list entry and an approved label were each searched for and none was located, so the approval is recited here and not established here. Do not write that CDSCO approved the drug on that date, which upgrades a recital into a ruling. Do not write the categorical negative either, because no instrument in this build supports it. For a patient, availability is checked against a current label or a pharmacy record, never inferred from a committee minute.

18.5 Why this one row licenses every declared gap in the India column

This row does more work than its own drug. It is the strongest of the three legs establishing that the CDSCO published approval-list corpus is incomplete, because here CDSCO contradicts itself: its own committee recites an approval that appears in none of the 66 published lists, and the list that should have carried it declares coverage of the date. Every other Indian entry in this chapter depends on that incompleteness, so one compact agent profile is why the rest of the column may be written as declared gaps at all.

The second leg is an era mismatch, and it bears on a chapter of recent approvals: every drug this corpus is known to find sits in a document dated 2023 or earlier, and independent sensitivity in the 2024 to 2026 window is 0 of 1, that one being this drug, whose Indian approval is established from outside the corpus and which the corpus misses. The ledger refuses to count the one agent here that the corpus does find as a second data point, because scoring a corpus hit as proof that the corpus finds hits is circular, and the positive-control tally of 17 of 20 may not warrant a gap on any agent approved after 2023.

848 of 848	0 of 2730	0 of 66	0 of 121
CHARACTERS, US ROW VERSUS LABEL	EU REGISTER MEDICINE ROWS MATCHING	CDSCO APPROVAL LISTS WITH AN ENTRY	EXTRACTED CDSCO TEXTS NAMING THE DRUG

Quote the conclusion rather than improving on it. Absence from this corpus is evidence and not proof; no India row in this build asserts that a drug is not approved in India; each asserts instead that no CDSCO approval record was located, and names what was searched. A categorical negative is a claim about a country, and the instrument only ever supported a claim about a search.

18.6 A date and the label on the date are one record

One correction inside this row is the most transferable thing on the item, and it was made by hand after independent adjudication. The field describing what kind of date 26 December 2024 is once carried the wording the grounded Indian rows carry, which means an entry in an approval list. Anything joining the date to that description would have emitted a false Indian approval

sentence, correctly formatted and entirely wrong. It was raised as a blocker twice by independent cross-family review, and the row's state was left unchanged, because partial on a supportive source was already correct. A date and the label describing it are one record and must be corrected as one unit, and a string true in one row is not therefore true in another.

The falsification pass makes the point mechanically, and its verdict on the Indian row is a term worth owning: **unresolved extraction**. All three text instruments returned it rather than a pass, because this row cites no approval-list entry and therefore offers no indication cell to read. The finding was recorded as measured rather than assumed, through the run-time scan of all 121 extracted texts and a second source check across five cached committee-minute files, in one of which the quoted sentence appears verbatim. The discipline is written into that file in one line: a failed extraction is unresolved extraction, never a pass. An instrument that could not read is not an instrument that found nothing.

MNEMONIC

Recited, not Ruled

Recited is what a committee minute does: it reports an approval. Ruled is what a permission, an approval-list entry or an approved label would do: it constitutes one. The Indian row holds the first and not the second, so the sentence it licenses names the minute, its date and its verb. Naming CDSCO as the actor on 26 December 2024 upgrades a recital into a ruling.

Reduced to three sentences, the item is this.

- The mechanism acts on serotonin, dopamine and glutamate at once, and is a rationale for the licence, never a substitute for reading it.
- The American record is grounded and quotable: schizophrenia in adults; bipolar I or II depression alone or added to lithium or valproate; major depressive disorder only as an addition to an antidepressant.
- The Indian record is a recited approval and not an established one, and that single row is the evidence that the published approval-list corpus is incomplete.

One molecule across three transmitter systems, one licence quoted whole and left inside its own country, and one Indian record that recites an approval no published list prints.

19 · Zuranolone as an individual agent profile (cross-reference to recent advances)

This is the prescribing half of zuranolone, and it is short on purpose. The mechanism, the class lineage, the postpartum against major depressive reading and the regulatory chronology belong to the zuranolone and neurosteroid antidepressants item in this chapter. What follows is the agent-reference residue: the dose, the pharmacokinetics, the interactions, and the points where two labels describe one molecule differently.

19.1 Identity, and the border running under every line below

Zuranolone is an oral neuroactive steroid, classified in the European document under psychoanaleptics, other antidepressants, with ATC code N06AX31, supplied as 20 mg, 25 mg and 30 mg hard capsules. Two regulatory records are held here and a third is empty. The United States label reads that ZURZUVAE is indicated for the treatment of postpartum depression (PPD) in adults, on Drugs@FDA application NDA217369, whose earliest approval-status date in the record is 4 August 2023, status true on 1 August 2026. In the European Union the medicine is Zurzuvae, EMEA/H/C/006488, holder Biogen Netherlands B.V., authorised throughout the Union from 17 September 2025; its indications section names one indication only, quoted verbatim in the neurosteroid item and deliberately not paraphrased here. For India, no CDSCO marketing approval record was located as of August 2026, on the authority of the Central Drugs Standard Control Organisation and the Drugs Controller General of India.

Everything below is label content, and a label describes what one regulator permitted inside its own territory. A dosing table is not a permission, and no figure here crosses a border.

19.2 One dose, one evening, fourteen days, and what moves it

SITUATION	DOSE	WHAT THE LABELS SAY
Standard	50 mg orally once daily in the evening for 14 days	The European document calls this a single course; the United States label states that safety and effectiveness beyond 14 days in a single treatment course have not been established
Not tolerated	40 mg once daily, same 14 days	European trigger is failure to tolerate 50 mg, United States trigger is depressant effects arising inside the period. The European document adds that dosing may be stopped without down-titration
Strong CYP3A inhibitor	**30 mg once daily through the 14 day treatment period**	Both. The United States label adds that no dosage modification is recommended with a moderate CYP3A4 inhibitor
Moderate or severe renal impairment	30 mg, same period	European: moderate eGFR 30 to 59 mL/min, severe under 30 not requiring dialysis. United States: moderate or severe, eGFR under 60 mL/min/1.73 m2. Mild, 60 to 89, unchanged in both
Severe hepatic impairment	30 mg, same period	**Child-Pugh class C** only; class A and class B unchanged in both
Missed evening dose	Nothing extra	Take the next dose at the regular evening time, take no additional capsules the same day, and continue to the end of the course. Both

The capsule arithmetic goes wrong on the prescription, not in the answer. Fifty milligrams is two 25 mg capsules, forty is two 20 mg capsules, thirty is a single 30 mg capsule, so every dose change is a change of strength, not of count.

The administration instruction is an exposure requirement. Capsules are swallowed whole in the evening with fat-containing food, as either a meal or a snack, quantified by the United States label as 400 to 1,000 calories and 25 to 50 per cent fat and left by the European document as everyday examples. Zuranolone may be given alone or alongside stable background oral antidepressant therapy, so it is neither an add-on by licence nor a replacement for one.

RULE **Every dose in this entry is a fourteen-day dose.** There is no titration upward, no loading and no maintenance limb. Both labels specify each reduced dose for the whole treatment period rather than for one evening, so a patient started at 30 mg for an inhibitor, a kidney or a Child-Pugh class C liver stays there to the end. The dose is decided once, at the prescription; the only trigger arising later is tolerability.

19.3 The pharmacokinetics that explain the instructions

Two of these numbers are why the patient is told anything at all. Each is printed with the document it came from, because the two labels do not always use the same measure.

PARAMETER	VALUE	DOCUMENT, AND ANY DIVERGENCE
Time to peak concentration	**5 to 6 hours after the dose**	Both
Steady state	Reached in 3 to 5 days, accumulation about 1.5-fold	Both
Distribution	Volume of distribution greater than 500 L; plasma protein binding greater than 99.5 per cent	Both; the European document adds that the volume is dose-independent
Terminal half-life	**Approximately 19.7 to 24.6 hours in an adult population**	Both
Mean apparent clearance	32.7 L/h European, 33 L/h United States	One quantity, two roundings; quote the one whose document you name
Excretion	45 per cent in urine and 41 per cent in faeces, both as metabolites, with negligible unchanged drug in urine and under 2 per cent in faeces	Both. Renal impairment matters through exposure, not unchanged excretion
Food effect at 30 mg	European: peak up 228 per cent and AUC up 55 per cent low-fat, 334 and 90 per cent high-fat, against fasting. United States: peak up about 3.5-fold and AUClast about 1.8-fold low-fat, 4.3-fold and 2-fold high-fat	Two documents, two metrics; time to peak unaffected in both

50 mg THE ONLY DAILY DOSE	**30 mg** THE REDUCED DOSE, THREE TRIGGERS	**5 to 6 h** TIME TO PEAK CONCENTRATION	**19.7 to 24.6 h** TERMINAL HALF-LIFE

A half-life near a day with once-daily dosing keeps the course at steady state throughout, so the twelve-hour restriction both labels attach to each dose is not a wash-out and does not lengthen as the fortnight goes on. It is set against a performance threshold, worked through in the neurosteroid item, not against elimination.

The exposure figures behind those thresholds: in Child-Pugh class C the peak concentration was 24 per cent lower and AUCinf 56 per cent higher than in matched healthy subjects, while class A and class B were unchanged, which is why only the severe class moves the dose. Exposure rose in moderate and severe renal impairment, and the drug has not been studied at an eGFR under 15 or in patients requiring dialysis. No adjustment is made for weight, race or age; Black or African American subjects showed a 14 per cent higher apparent clearance, which the European document itself calls not clinically meaningful. There is no geriatric experience in this indication.

19.4 Interactions: a victim of CYP3A, almost never a perpetrator

Metabolism is extensive with CYP3A primary, so the interactions worth knowing run one way. A **strong CYP3A inhibitor** raises exposure, quantified as AUCinf up 62 per cent with itraconazole, and the answer is 30 mg rather than avoidance; the European examples are protease inhibitors, azole antifungals and some macrolides, with grapefruit avoided for the same reason. Induction reverses both magnitude and management: AUCinf falls 85 per cent with rifampin, and inducers are avoided outright with no dose-adjusted alternative. The European list names carbamazepine, phenobarbitol, phenytoin, primidone, rifampicin, St John's Wort and efavirenz; the last three are the ones candidates omit.

In the other direction the drug is quiet. Repeated zuranolone did not alter exposure to simvastatin, a CYP3A substrate, or bupropion, a CYP2B6 substrate, and it is not a substrate of P-glycoprotein, nor expected to inhibit the common cytochromes or the major transporters at clinically relevant concentrations. The one important interaction that is not pharmacokinetic at all is with other central nervous system depressants: repeated 50 mg daily doses with alcohol or alprazolam produced increased impairment of psychomotor performance, yet the United States pharmacology records that alprazolam and ethanol did not alter zuranolone exposure, and zuranolone did not alter theirs. The additive effect is pharmacodynamic, so no concentration would predict it, and the management is dose reduction where the combination is unavoidable.

19.5 Where the two labels do not say the same thing

Contraindications first, because the divergence is total. The European document contraindicates hypersensitivity to the active substance or excipients, and pregnancy. The United States contraindications section reads, in its entirety, None. That is not a disagreement about foetal risk: both advise effective contraception during treatment and for one week after the final dose, and the United States label carries an embryo-fetal toxicity warning on the same animal findings. One regulator used a contraindication, the other a warning.

Dependence next. The European special-warnings section states that, on clinical-trial data, zuranolone has low physical dependence potential. The United States label states that it may produce physical dependence, and prints the evidence: in healthy subjects given 50 mg for 6 to 7 days, discontinuation was followed by insomnia, palpitations, decreased appetite, nightmare, nausea, hyperhidrosis and paranoia, all mild to moderate. In that jurisdiction zuranolone is a **Schedule IV controlled substance**, a classification belonging to United States law that must not be carried elsewhere. Quoting either sentence and calling it the position on **physical dependence** quotes half a field.

The United States label also carries the antidepressant-class warning on suicidal thoughts and behaviour, with pooled figures showing 14 additional patients per 1,000 under 18 and 5 additional at 18 to 24 against 1 fewer at 25 to 64, adding that zuranolone does not directly affect monoaminergic systems. The European special-warnings section, read whole, carries five headings, on driving, depressant effects, interactions, abuse and dependence, and excipients; suicidality is not among them.

The paediatric sentences differ in kind and must not be merged. The European document records that safety and efficacy in postpubertal females less than 18 years old have not been established, and that there is no relevant use in prepubertal females. The United States label records only that safety and effectiveness in paediatric patients have not been established. Reading either as a gloss on the other manufactures a floor neither document sets.

TRAP **Quoting a frequency without its population.** The European pooled figure for somnolence among subjects given 50 mg is 26.5 per cent; the United States table for the same dose in its Study 1 gives 36 per cent against 6 per cent on placebo. Both are correct and both are 50 mg figures; the poolings differ. The trap runs the other way on tolerability: the European frequency for dose reduction or interruption is 14.3 per cent, while the United States figure for reduction alone is 16 per cent, so the narrower construct carries the larger number, which is possible only because the denominators are different patients.

19.6 The prescribing points, in the order they are decided

One intentional overdose was reported before marketing: 330 mg, about 6.5 times the maximum recommended dose, producing an altered state of consciousness that resolved the next day after intravenous fluids. There is no specific antidote and management is supportive.

- Decide the dose before the first capsule: a strong CYP3A inhibitor, a moderate or severe eGFR, or a Child-Pugh class C liver each fixes the course at 30 mg, and all three are readable from the chart.
- Check inducers separately, because they are an avoidance and not a dose problem: carbamazepine, phenytoin, phenobarbitol, primidone, rifampicin, St John's Wort, efavirenz.
- Counsel the evening, the fat and the twelve hours as one instruction with three parts.
- Name the jurisdiction whenever a status, a schedule, a contraindication or a frequency is quoted, because each of those four differs between the documents or exists in only one.

GOOD TO REMEMBER

Three of the four reasons to reduce are known before the prescription is written and all three land on the same 30 mg capsule; the fourth, tolerability, lands on 40 mg and arises during the course. Settle renal function, liver status and the interacting list at prescribing, and keep the 40 mg step in reserve for the first two evenings.

IN INDIA

Every figure in this entry comes from two foreign labels, and a label describes a permission granted elsewhere. As of August 2026 no CDSCO marketing approval record was located for zuranolone; the missing document is an affirmative CDSCO or DCGI new-drug permission, an approved Indian label, or an entry in an official approval list. A CTRI registration, expert-committee minutes and a brand in a distributor catalogue are none of them an Indian marketing authorisation. This item holds no instrument that measured Indian availability in either direction, and the measured limits of the approval-list corpus are set out in the neurosteroid item.

MNEMONIC

Evening, Fat, Fourteen, Thirty

Evening is when it is given and why the restriction runs into the next morning. Fat is the condition of absorption. Fourteen is the whole exposure, with no titration and no second course. Thirty is the reduced dose and its three triggers, all knowable before the first capsule.

One dose for fourteen evenings, one reduced dose for three chart-readable reasons, and a jurisdiction attached to every status, schedule, contraindication and frequency you quote.

Consolidate and apply

20 · Worked vignettes

A vignette in this chapter is almost never a question about a drug; it is a question about a record. The stems that carry marks open with somebody wanting something: a patient who has read about a new treatment, a colleague who prescribed it in another country, a committee asked to stock it. Then the stem plants the fact that decides the answer, and in this territory that fact is almost always a jurisdiction, a document, a date, or a kind of permission. What the closing vignettes rehearse is the reading skill the chapter was built to install: finding the status question inside a clinical stem, and answering it at the grain the evidence actually supports.

The discipline that earns marks is refusing to answer a regulatory question with a pharmacological one, and refusing to answer it about the wrong country. A patient asking whether they can have a treatment they have read about is not asking what it does at the receptor; they are asking whether a named regulator has authorised it for them, here, now. A candidate who supplies the American record for a patient in an Indian outpatient clinic has answered the right question in the wrong place, which costs more because it looks like knowledge. Where a stem turns on mechanism, or on choosing within a class, that reasoning belongs to the Mood disorders: pharmacotherapy notes and the Psychopharmacology I: principles and core drug classes notes. This chapter supplies the layer above the prescription: whether the thing may be prescribed at all, on whose authority, and how to say so when the answer is not known.

20.1 Four questions locate the status claim

Before deciding what a recent-advances stem is asking, run the same four questions across it. Between them they cover almost every way an examiner can turn a request for a new treatment into a regulatory problem, and each opens onto a different part of this chapter.

Border	Document	Date	Verb
WHICH JURISDICTION IS ASKED	WHAT WOULD SETTLE IT	WHEN IT WAS TRUE	WHAT KIND OF PERMISSION

Border asks which country's record the stem is asking about, and that is set by where the patient is, not by where the literature was written. Document asks what record would settle the question and which body issues it: a published approval list, a current approved label, a public assessment report, a commission decision. Date asks when the status was last checked, because a regulatory statement with no date is a claim about the present made from an unspecified past. Verb asks what kind of permission is described, since permission to run a trial, an approval-list entry, a recital inside committee minutes and a product being supplied are four different facts about one agent. A well-built stem makes one of the four decisive and leaves the others as scenery.

RULE **A status sentence with no jurisdiction, no source and no date is not an answer.** Those three fields, with the indication itself, are what a regulatory claim here has to carry, and a sentence missing any of them cannot be checked. "It is approved" is not a shorter version of the answer; it is a different and unsupportable claim, because it silently asserts the same status everywhere. If you cannot give all three, say what was searched and what would settle it: a complete answer of a different shape, not a weaker version of the same one.

20.2 The vignette that resolves on a border

The commonest stem names an agent and asks whether the patient in front of you can have it, and the candidate answers from whichever country's record they happened to revise. The reading move is to fix the border first, and to treat the answer as a property of the pair, agent and jurisdiction, rather than of the agent. Regulatory maturity does not run in one fixed order, and this chapter holds the counter-example. One of the orexin receptor antagonists holds an Indian approval-list entry with an insomnia indication, dated **1 July 2021** under the printed column heading Date of issue, while for the same substance no European centralised procedure of any kind was found on record, both status true as of 1 August 2026. A candidate who reasons that an agent must be available in India because it is available in the United States, or that it cannot be because Europe has no record of it, has assumed a ladder that does not exist.

The other half of the move is knowing what a European absence is worth. The instrument behind it is the agency's own register of centrally authorised medicines, which records refusals, withdrawals, lapses, revocations and suspensions alongside authorisations, so an agent absent from it has never concluded a central application of any kind. That is a real finding, and narrower than it sounds, because the register covers the centralised procedure only and a product authorised nationally through the decentralised or mutual-recognition route would never appear in it. One further weakness belongs in a careful answer: the confirming second instrument ran on

all eight grounded rows and none of the eight absences, so no European absence here has independent corroboration.

20.3 The vignette in which the record was searched and not found

A second family hands the candidate an agent for which nothing was found, and the examinable content is entirely how that is said. A patient asks for an intranasal treatment they have read about; the Indian record was searched, and no approval was located. The wrong answer converts an absence of record into a regulator's negative finding. The right answer is a **declared gap**: name what was searched, name the authority that would settle it, and name the document that is missing. For the Indian column that means three instruments, namely a corpus of published approval lists, a search of the regulator's own domain in which every returned document is read rather than counted, and a search of the national trial registry run with a positive control. **Marketing authorisation** is what would settle it, and nothing softer will do.

The instrument's blindness has been measured, which is why this is not pedantry. The published approval lists are demonstrably incomplete, shown most strongly by the regulator contradicting itself: its own expert committee minutes of April 2025 recite an approval for an agent that appears in **0 of 66** of those documents, and the list declaring calendar-2024 coverage states its own range in its running header, so it is not truncated before the recited date. A corpus that misses a real approval cannot then be offered as evidence that some other approval does not exist. Absence here is evidence and not proof, and the sentence to write is that no approval record was located, naming what was searched.

IN INDIA

Write it in this order. The authority is the Central Drugs Standard Control Organisation and the Drugs Controller General of India. Expert committee minutes, **clinical-trial no-objection certificates**, International Cell written confirmations, CTRI registration, a brand listing, pharmacy availability and manufacturer copy are none of them marketing authorisation. So the honest sentence to a patient or a referring colleague is that no Indian approval record was located as of the date you checked, that what would close the question is a CDSCO or DCGI new-drug permission, an approval-list entry or a current approved Indian label, and that a pharmacy being able to source something is a different fact.

TRAP **Converting a search that found nothing into a regulator's negative finding.** Writing that an agent is not approved in India, or not approved in the United States, asserts that a body has stated something no instrument here asked it. The probes measured named routes and named documents; they are silent on what a regulator would say if asked directly. Worse, an American gap in this build has already been wrong once: one agent's record carries no cross-reference block at all, so three of five routes returned not-found on it while the active-ingredient route returned the record. A declared absence is written at the strength the instrument supports and never one notch harder, and an examiner marks the notch.

20.4 The vignette in which the approval stands and the product does not

A stem describes a woman with severe postpartum depression whose team has read about an intravenous neurosteroid and wants to know whether it can be arranged. An approval and a marketing status are two separate facts about one product. The American record for the first neurosteroid in this class carries an approval dated 17 June 2019 together with marketing status Discontinued. Both fields, approval and Discontinued marketing status, sit on that one record, which carries nothing relating them, so an answer gives both and reasons from neither to the other. Naming the approval and dropping the Discontinued marketing status tells the family half of what the record says. Saying the drug was withdrawn as unsafe is the more instructive overreach, because it reads a reason into a field that gives none: the Discontinued marketing status is a fact about the product rather than about the approval, and no field on this row states why marketing stopped or names any action taken about it by a regulator or a sponsor, so a candidate who supplies one has left the record, in either direction. Say that the approval stands on the record and the marketing status is Discontinued, and stop where the record stops.

The same row, the Discontinued one, carries a second object worth a sentence in a viva. Its date of approval and the date its approval letter was signed are different quantities, because the approval letter sets the date of approval as the date of a later scheduling notice, and the earlier date has been retracted from the field a reader would take a first-approval date from. Publishing a signature date under a first-approval field name is a true date of one thing wearing the field name of another, which is the commonest way a correct number becomes a false claim. Where an answer usually lands instead is the oral agent in the same class, whose rows sit in this chapter's neurosteroid section.

20.5 The vignette that resolves on the kind of document

Four different things get called approval in conversation and only one is a marketing authorisation, so a stem can be built entirely out of the difference. The first is a trial permission. The Indian file for the muscarinic combination is documented at pre-approval stage: no-objection certificates name the combination, and committee minutes for neurology and psychiatry dated 30 May 2025, 10 December 2025 and 22 December 2025 record a named sponsor presenting Phase III protocols, with the committee recommending permission to conduct the trial. That is affirmative evidence about the file and it is not an approval. The second is a recital: minutes of a **subject expert committee** mentioning an approval for another agent license the statement that those minutes recite such an approval, and do not license the statement that the regulator approved the drug on that date. The third is supply, a fact about a market. The fourth is a strength printed beside a product name in a list, which is part of the product's identity and not a dosing instruction.

The value sits in the three errors this catches. A candidate who writes that the muscarinic combination is approved in India has upgraded a trial permission into a marketing authorisation. One who writes that the other agent was approved on the recited date has upgraded a recital into the permission itself, when the primary permission or label was never located. And one who writes that either agent is unapproved has downgraded affirmative documentary evidence into a

negative nobody established. A committee recommending that permission be granted is a step towards permission and not the permission, so say which document you are describing before you say what it shows.

20.6 The vignette that resolves on a record that is authoritative and wrong

The next family teaches a reading habit rather than a fact. A record can come from the right regulator, cite the right document and still say something false, and both shapes it takes are silent. The first is truncation. The European indication for the oxybate medicine reads Treatment of narcolepsy with cataplexy in adult patients, adolescents and children from the age of 7 years. An earlier stored version of the same field read "Treatment of narcolepsy with cataplexy in adult patients." The words "adolescents and children from the age of 7 years" had been dropped and the surviving prefix closed with a full stop the source does not have. Nothing looked wrong, because a cut closing with valid punctuation is invisible to the eye. The vignette that follows is paediatric: a child with narcolepsy and cataplexy, and a candidate whose remembered indication has quietly excluded them.

The second shape is a wrong product wearing the right citation. The cited application and the date were both correct and both belonged to one product; the indication text was not, because it was the text of a different application by a different sponsor, whose label reads almost identically in structure. A reader comparing structure or length would see nothing. Both defects were found the same way, by somebody reading a stored field, noticing it did not close properly or did not name the product it claimed to, and refusing to complete it from memory. That refusal is the transferable skill: check the product's identity against the application's, quote an indication whole or say you are paraphrasing, and treat a quotation ending on a tidy full stop as unverified.

20.7 The vignette that resolves on what the evidence does not license

Not every stem here is regulatory. The other family gives a finding and asks what follows, and the mark sits in the inferences refused rather than the ones drawn. The umbrella review of the serotonin literature is the standing example. What it examined and what it reported are one question; what it licenses is another; and it does not license the conclusion that antidepressants do not work, that depression is not biological, that serotonin has no role in mood, or that evidence about aetiology settles a question about efficacy. Aetiology and efficacy are separate claims answered by separate designs, and a stem offering the review as grounds for stopping a patient's treatment is testing exactly that separation.

The same discipline runs through the newer research territory. A psychedelic-assisted trial is an intervention package, drug with preparation, setting and psychological support, so a question about what worked is a question about which component the design can separate, with expectancy, functional unblinding, therapist effects and the exclusion of higher-risk patients all on the appraisal path; and an access arrangement for a named patient is not a product approval. A stem describing an algorithm, a candidate biomarker or a prescription application is asking which rung the evidence has reached, across analytical validity, clinical validity, clinical utility, regulatory authorisation and routine deployment. No rung is inherited from the one below, and a **measured absence** of evidence at one rung is not a claim about the rung beneath. Those

arguments are developed in this chapter's serotonin-hypothesis, psychedelic-therapy, machine-learning, biomarker and digital-therapeutics sections; an answer here names the rung and stops.

20.8 Look-alike stems and the answering spine

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

TWO STEMS THAT LOOK ALIKE	THE FACT THAT SEPARATES THEM	THE AXIS IT LANDS ON
An agent available in one country, against the same agent asked about in another	Which jurisdiction's document was actually read, since the answer belongs to the pair and not to the drug	Border
An absence of record, against a regulator's negative finding	Whether any body has stated that nothing exists, or only that named searches found nothing	Document
An indication quoted whole, against one truncated at a full stop the source does not have	Whether the population at the end of the sentence survived the quotation	Document
A clinical-trial permission, against a marketing approval	What the document permits: conducting a study, or marketing for an indication	Verb
Committee minutes reciting an approval, against the approval record itself	Whether the permission, the list entry or the approved label was ever located	Verb
A status checked this month, against one carried in from memory	Whether the claim carries the date on which it was true	Date

With the axis identified, an answer that earns marks follows four moves in order.

- Fix the border. Name the jurisdiction the question is being asked in, before naming any status at all.
- Name the document that would settle it, and the body that issues that document.
- Give the verb its correct strength: authorised, permitted for a trial, recited in minutes, supplied, discontinued.
- Say what the record does not settle, and what you would have to read to close it.

That fourth move is where regulatory reasoning becomes visible to an examiner. A candidate who answers that a drug is approved has made a claim about a molecule; one who answers that a named regulator authorised a named product for a named indication, on a document they can identify, true as of a date they can give, and that the Indian record was searched on three named instruments with no approval located, has made a claim about a record, and a claim about a record can be checked. Every vignette in this territory rewards the second answer.

Answer at the grain of agent and jurisdiction, with the document, the date and the strength of the verb attached, and turn a search that found nothing into a declared gap rather than into a negative.

21 · Consolidation and quick review

This page teaches nothing and asks everything. The eighteen sections behind it have each been taught once, in the section that owns them, and consolidating them is not rereading them. It is finding out what can be produced from a blank page with the book shut, because that is the only form in which any of this will be available in a viva. What follows is a list of prompts rather than a summary: each names a section and the answer it should be possible to build cold.

Two habits make it easy to consolidate badly. One is revising the chapter as eighteen unrelated topics, which yields eighteen lists and no method. The other is holding a regulatory position as a property of a molecule; recalled that way it is unmarkable rather than merely weak, and it is usually wrong about which regulator it belongs to.

21.1 The two things to hold before any prompt below

The first is the unit of recall. Nothing here is held at the level of an agent: a status is held at the level of the **agent and jurisdiction cell**, and the cells of one agent are independently sourced and free to disagree. The ledger behind these notes carries **48** such cells, and this chapter holds pairs in which one cell is grounded while its neighbour is a declared gap. A single verdict about an agent is therefore not a compressed answer but a claim about cells that were never read. Producing one is the **border crossing** this chapter exists to prevent, and it costs marks both ways: an approval reported for a regulator holding no record, and a categorical negative written where only an instrument was silent.

The second is the audit. Four questions to run against any status answer the moment it is written, one per field such a sentence carries, the last being the date the cell was true, which these notes record as status true as of.

- Which regulator is this sentence about, and does the sentence itself say so?
- Where did the wording come from, and is it the document's own wording rather than a paraphrase that has quietly tightened it?
- On what date was this true, and does that date appear beside the claim?
- If the answer is a negative, what was actually searched, and whose authority would the claim finally rest on?

RULE **Recall by cell, never by agent, and produce the sentence rather than the verdict.** An answer naming an agent and a status has answered a question nobody can mark. An answer naming the regulator, the wording, the document and the date has answered the question asked, even when what it reports is a gap.

18	48	1	0
SECTIONS TO PRODUCE COLD	AGENT AND JURISDICTION CELLS BEHIND THEM	CELL A STATUS SENTENCE MAY SPEAK FOR	EVIDENCE LANES THIS CHAPTER HAS RUN

21.2 Twelve prompts for the recent advances band

One prompt per section. The middle column is what to produce; the right-hand column is what an otherwise fluent answer still misses, which is where the marks sit. No cell supplies the answer it asks for, with two exceptions the chapter's own rules force: a declared gap can be named only by naming what was searched, and one product's approval can be named only beside its marketing status, which a hard rule of this chapter requires in the same sentence as the product.

SECTION	PRODUCE FROM A BLANK PAGE	STILL INCOMPLETE WITHOUT
Esketamine and intranasal therapy	The indication each regulator's own record carries in its own words, the monotherapy and adjunctive positions separated per regulator, and any limitations of use the United States wording attaches	A jurisdiction on every clause, and the Indian cell written in whichever of the three states its own record supports
Zuranolone and the neurosteroid antidepressants	The neurosteroid rationale, the postpartum and the major-depression evidence read separately, and what brexanolone teaches about an approval that stands while the product's United States marketing status is recorded as Discontinued	The two populations kept apart, and the reason a date of approval can be set by something other than the approval letter
Rapid-acting antidepressants as a class	What makes the category a category rather than a list, and its members placed against onset, durability, evidence maturity, treatment setting and regulatory maturity	Each axis ordered qualitatively with its basis named, and not one figure on any of the five
Dual orexin receptor antagonists and sodium oxybate	The agents cell by cell, and the pair of cells that defeats reading the three regulators as a ladder of maturity	Whichever population a quotation of the European oxybate wording silently drops when it is cut at a full stop
Xanomeline, trospium and the muscarinic antipsychotics	Why a peripheral antagonist is paired with the agonist, and where the strategy sits beside rather than inside the dopamine algorithm	Which of the two named substances the Indian approval-list corpus cannot display, and what that does to any negative about the pair
Endoxifen and the protein-kinase-C-targeted agents	The Indian cell and the United States cell as two sentences that never meet, with protein kinase C inhibition named as a rationale and not a result	The refusal to merge the two cells into one agent-level statement; the United States cell given as the named routes that were run against Drugs@FDA and found nothing, never as a verdict, and any negative left to rest on the Food and Drug Administration's own reading of that register; and no comparison against an established mood stabiliser
Anti-amyloid monoclonal antibodies	The class of claim these agents exemplify, why this chapter prints no position for them, and what an Indian old-age service would have to establish first	The delegation named as one, the material handed to the Dementias and neurocognitive disorders notes, and no status and no date anywhere

SECTION	PRODUCE FROM A BLANK PAGE	STILL INCOMPLETE WITHOUT
The serotonin hypothesis debate	What the umbrella review examined and what it reported, then the four inferences it does not license	All four stated, and evidence about aetiology kept apart from evidence about efficacy
Psychedelic-assisted therapy	The intervention as a package of drug, preparation, setting and psychological support, with psilocybin and MDMA both named and neither given a status	Expectancy, functional unblinding, therapist effects and the exclusion of high-risk patients marked on the appraisal, and access arrangements kept apart from product approval
Artificial intelligence and machine learning	What a performance claim is bound to before it can be read at all, and where external validation sits in that list	The step from performing to deployable stated as a step, and no performance figure of any kind
Biomarker-driven and computational psychiatry	Analytical validity, clinical validity and clinical utility separated, with the evidence each one actually needs	The rungs kept in order, and the reason a marker result licenses no conclusion about a patient
Digital therapeutics and prescription apps	What a prescription digital therapeutic is as a regulatory object, why it is a device and not a drug, and what engagement and attrition do to an effect that a tablet does not have	The blindness of this chapter's instruments stated on your own authority, no regulatory verb on any named product, and deployment left to the Community psychiatry and public health notes

21.3 Six prompts for the individual agents

These sections are compact and are marked on the same two things: what the receptor signature commits the agent to, and what each regulator's record says. Write each as cells and a pharmacology, not a paragraph about a drug.

- Asenapine: the receptor signature, the formulation constraint that follows from it, and then, cell by cell, which regulators hold a grounded record and what each one's wording covers.
- Lurasidone: the metabolic and movement profile and the indications, then the reason no first Indian approval date appears in this chapter, and what the Indian corpus does hold for it instead.
- Brexpiprazole: the receptor signature and the indications it carries, and what its United States record teaches about a citation whose fields do not all point at the same thing.
- Cariprazine: the receptor signature that separates it within its own class, the indications, and what a verbatim quotation loses when it is cut at a fixed character count instead of at a clause.
- Lumateperone: the receptor signature and the indications, then the jurisdiction whose record is supportive-source-only rather than grounded, and the sentence that record does not license.
- Zuranolone as an agent profile: the compact agent fields alone, the neurosteroid case left where it was argued, and what a profile may carry that a narrative section already owns.

21.4 A second pass, sorted by the way an answer fails

Once the eighteen prompts run cleanly, run the chapter in the other direction, sorted by how an answer goes wrong rather than by topic, because a supplementary question arrives as a defect and not as a heading.

HOW THE ANSWER FAILS	THE PROMPT THAT TRAINS IT	WHERE IT IS TESTED
A status carried across a border	Take any agent this chapter names and write its cells one at a time, refusing to let a full cell fill an empty one	The esketamine, orexin and endoxifen sections, and the agent prompts above
An answer given at the wrong grain	For every agent whose cells do not read alike, say why a single agent-level verdict cannot be written at all	The orexin and endoxifen sections
A negative built out of a silence	For each gap, say which instrument was pointed at it, how sensitively, and whose authority a claim would finally rest on	The muscarinic, psychedelic and digital-therapeutic sections
A recited approval read as a granted one	Name the state that is neither a grounding nor a gap, and the sentence it licenses in place of the obvious one	The lumateperone section
A figure supplied where no lane holds one	For each evidence-driven section, list the quantities that are unavailable and say why writing one would be a fabricated citation	The neurosteroid, rapid-acting, serotonin-hypothesis, psychedelic, machine-learning and biomarker sections

21.5 Where these marks are actually lost

Three silences in this chapter look identical on the page and license entirely different sentences: one where a named instrument was pointed at the question and found nothing, one where the instrument could not have answered whatever the truth was, which is **instrument blindness**, so its quiet carries no information, and one where another chapter holds the fact and this one declines to restate it so two accounts cannot drift apart, which is a **delegated gap** and the only one of the three that is not an absence of fact. Say which of the three a silence is before writing anything from it; the section owning the third case works the discrimination out in full.

The other loss is quantitative. Every evidence lane behind this chapter is unrun, so no efficacy figure, effect size, response rate, trial name, sample size, endpoint, onset time, half-life, dose or adverse-event rate is available to any section in it. Supplying one from memory is not a weak answer but a **fabricated citation**, which nothing said afterwards can repair. Rehearse the substitution: say what is known qualitatively, name the ordering, and name what would have to be read for a quantity to be stated.

TRAP **Mistaking recognition for recall.** A section that reads as familiar on a second pass has been recognised, not retrieved, and recognition collapses under the first supplementary question. A prompt that withholds its answer leaves nothing to recognise. If a prompt cannot be written out in full sentences, with a jurisdiction on every status clause and a named source beside it, it is not yet known, however familiar the section felt.

21.6 The Indian limb, rehearsed as a separate answer

Write the Indian part of any answer separately, and last, because it is the limb most often filled in from whichever record the candidate has read.

IN INDIA

Write the Indian limb from an Indian document or do not write it. The authority is CDSCO and the Drugs Controller General of India, and a committee minute, a clinical-trial no-objection certificate, a CTRI registration, a brand on a pharmacy shelf and a manufacturer's own copy are none of them a marketing authorisation. Where no Indian record was located, the answer that earns the marks names what was searched, on what instrument, and whose authority a claim would rest on, and stops there rather than hardening into a categorical negative, because the corpus behind these notes carries a measured limit and an absence in it is evidence and not proof. Then take the question the examiner is actually asking: what an Indian service must establish first, meaning who confirms the diagnosis, who owns the decision, who carries the follow-up, and what the patient can afford.

21.7 What to carry into the paper

Consolidation buys faster access rather than more material, so the last week on this chapter should look nothing like the first. Four habits, in order.

- Produce rather than review. Write one section out cold each sitting and mark it against its source, instead of rereading three and feeling three times as prepared.
- Write the cells before the pharmacology. If the status will not come, no amount of receptor detail rescues the answer.
- Rework each vignette with the sections shut, because a vignette is where a border crossing surfaces as a decision rather than as a sentence, which is the form the examiner will use.
- Close every answer by naming what you hand to another chapter, or, where this chapter owns the section in full, by saying so and giving the full account.

What survives when the detail goes is short. Most questions here are about a record, so name the regulator before the agent; where a section has no regulatory record, name the construct and its limit instead. Every silence is a statement about an instrument, so say which instrument and how sensitively it looked. Every quantity is unavailable, so give the ordering and name what would have to be read. Every boundary is a sentence handed over rather than material rebuilt. Those four moves let a candidate answer on a section they revised badly.

Consolidating this chapter means producing it and not rereading it: one section at a time, cold, written as cells before pharmacology, with every status clause carrying its regulator and its source, and every silence carrying what was searched.

Index

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

#

5-ht2a 81
5-ht7 antagonism 119

A

abbreviated application 113
active metabolite 60
active placebo 80
acute monotherapy 112
acute suicidal ideation or behavior 19
acute tryptophan depletion 71
adjunctive therapy 112
adjunctive therapy to antidepressants 112
adjunctive therapy with lithium or valproate 112
adjunctive treatment 112
agent and jurisdiction cell 153
akathisia 112
akathisia rate 120
allopregnanolone 27
amstar-2 71
analytical validity 15
anti-amyloid monoclonal antibodies 15
apparent clearance 144
attrition 89
atypical antipsychotic 112

B

biomarker 13
border crossing 153
bulk drug 61

C

calibration 89
centralised procedure 19
change control 106
chemical imbalance 71
child-pugh class c 143
clinical utility 15
clinical validity 15
clinical-trial no-objection certificates 49
corrective model 75
ctri registration 24

D

d3-preferring 112
data leakage 90
dataset shift 89
date of approval 27
date of issue 48
declared gap 12
delegated gap 67
digital placebo 109
discrimination 88
disease-modifying 66
dual orexin receptor antagonists 15

E

earliest approval-status date in the record 21
engagement 67
era mismatch 43
expectancy 79
external validation 89

F

fabricated citation 12
four-field status claim 11
free-text search 64
functional unblinding 79

G

generative model 102

I

innovator application 113
instrument blindness 67
intended use 89
intervention package 81
intrinsic activity 112

L

large language model 89
limitation of use 112
limitations of use 19

M

maintenance monotherapy 112
major depressive episode 27

marketing approval 19
marketing authorisation 14
marketing authorisation holder 19
marketing status 26
measured absence 12
metabolised mainly via cyp3a4 119
minimally important difference 67
mixed features 61
monotherapy 18
muscarinic agonist 51
muscarinic antagonist 51

N

neuroactive steroid 26
no cdsco approval record was located 24
non-competitive nmda receptor antagonist 37

O

old-age psychiatry 66
optimism 89
orexin receptor antagonist 15

P

parameter recovery 102
partial agonist 112
physical dependence 32
positive allosteric modulation 27
positive control 14
positive controls 14
predictive 88
prescription digital therapeutic 89
protein kinase c 61
proxy label 89
psychiatric emergency 19

R

rapid-acting antidepressant 15
reference standard 89
regulatory recency 12
relative infant dose 32
release capacity 71
risk evaluation and mitigation strategy 27

S

schedule iv controlled substance 145
serotonin transporter 71
sham comparator 109
single course of treatment 27
software as a medical device 106
spectrum bias 89
status true as of 11
strong cyp3a inhibitor 29
structured product label 19
subject expert committee 19
subject expert committee minutes 19
sublingual tablet 112
summary of product characteristics 19
supervised learning 89
supportive source 12
surrogate marker 67
symptomatic 66
synaptic and extrasynaptic gaba-a receptors 28

T

target engagement 67
terminal half-life 144
therapist effects 79
tonic inhibition 27
transdermal system 112
treatment of schizophrenia in adults 53

U

umbrella review 69
unresolved extraction 42

W

written confirmation 24

About the authors

Dr. Wilfred D'souza

MD Psychiatry (Resident) · Certified Gestalt Therapist · Weave, Centre for Integrative Psychiatry

Works in integrative psychiatry with Schema Therapy as the primary lens, alongside CBT, DBT, motivational interviewing and emotion-focused therapy, and is a certified Gestalt therapist. His clinical interests run through neurodivergence, ADHD, bipolar disorder, personality disorders, addiction and developmental psychiatry: work grounded in lived experience of attention difficulties long labelled as laziness, which made psychiatry personal before it became professional. He has been developing the Toroidal Model of Psychosis and Affect (TMoP), his computational model of mind, and TMoP thinking shaped the Claude Code AI system he builds with: clinical software, teaching tools and much of this series' production pipeline among its work, bringing psychiatry into the AI era in his own way. Lead author of the Weave Sprint series.

Dr. Niharika Reddy

Consultant Psychiatrist · Weave, Centre for Integrative Psychiatry

Consultant Psychiatrist at Weave, Mumbai, seeing people online through weave.clinic. Her work centres on relationships and family dynamics, family therapy and women's mental health, with a particular focus on neurodivergence in women and on adolescent mental health. Practises integrative, person-centred psychiatry. Co-author of this volume.

THE WEAVE LIBRARY

You revised from one. The whole shelf is open.

Field guides and references written in practice, for the people who do the work and the people living through it. Most are free to read and download.



field guides · references · for clinicians and families · free to read

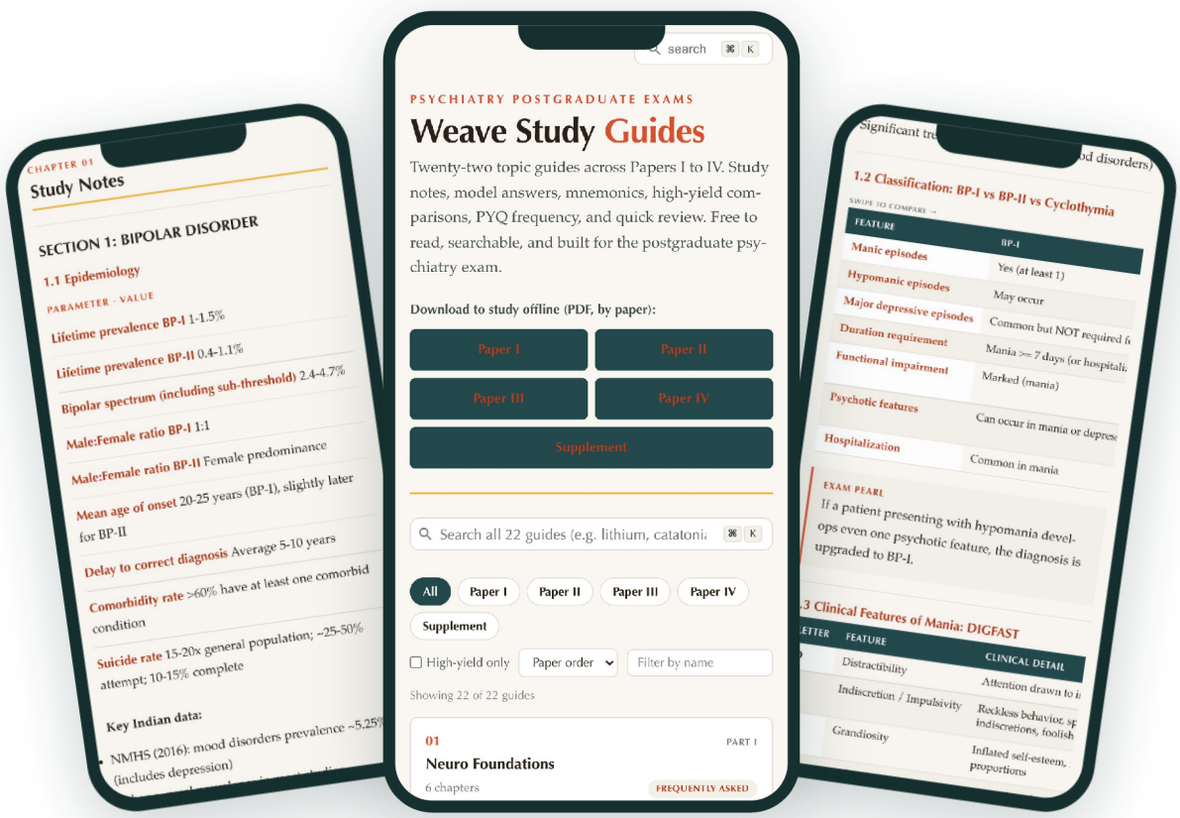
weave.clinic/library

TAP OR SCAN TO OPEN

WEAVE STUDY GUIDES

When the night is short, read the guide.

Twenty-two topic guides across all four papers. Study notes, model answers, mnemonics, high-yield comparisons and previous-question frequency. Free, searchable, built for this exam.



22 guides · all four papers · fully searchable · free to read

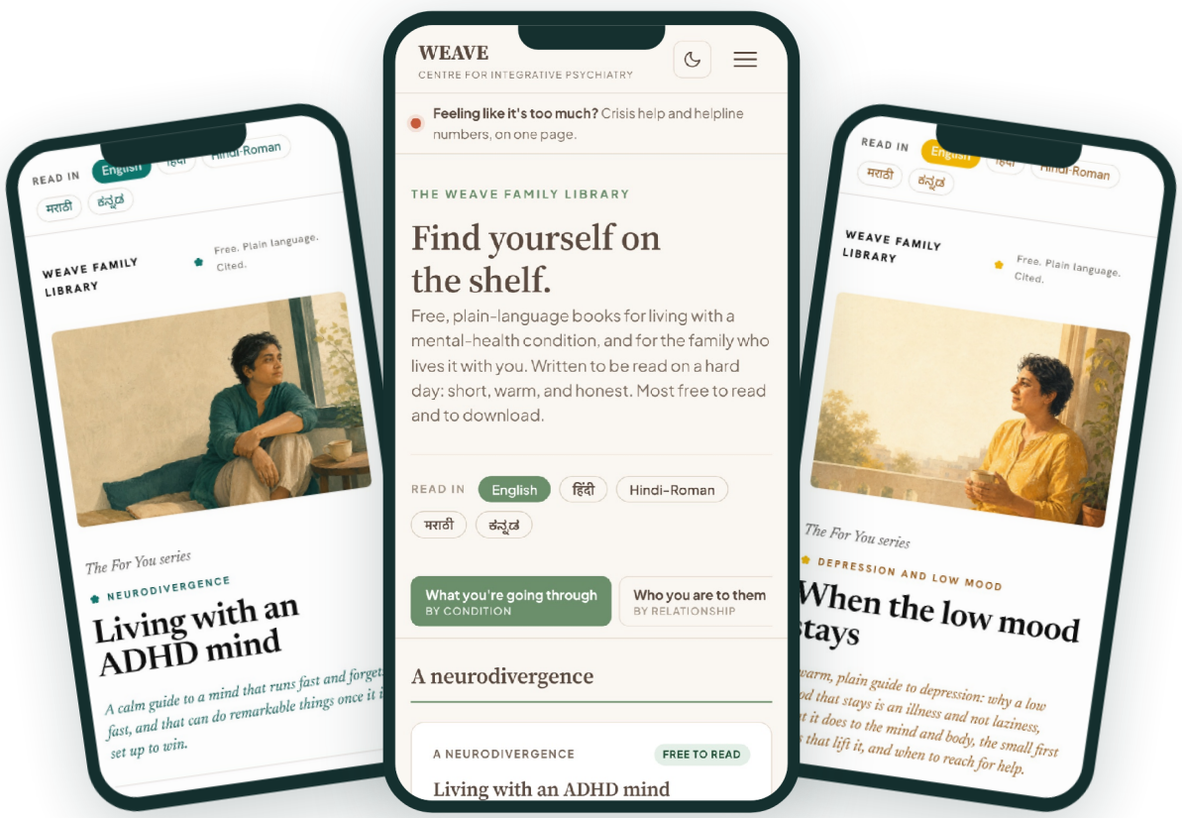
weave.clinic/study

TAP OR SCAN TO OPEN

THE WEAVE FAMILY LIBRARY

Send the family home with more than a prescription.

Free, plain-language books for living with a mental-health condition, written for patients and the people who love them. Find yours by condition, by relationship, or by situation.



plain language · patients and families · five languages · free to read

weave.clinic/family

TAP OR SCAN TO OPEN

A PAPER A DAY

The whole written exam, in one day's reach.

Every paper of the MD psychiatry written exam in one free book, with a colour memory code and clickable navigation, for the last week, the last night, the last hour before the hall.



four papers · one volume · colour-coded · free to read

weave.clinic/revision

TAP OR SCAN TO OPEN

STUDY TOGETHER

You do not have to revise alone.

These notes started as messages in a room full of residents helping each other through. That room is still open. Bring your doubts, your corrections, and the questions the textbooks do not answer.

Join the study group

https://chat.whatsapp.com/H7TsLGK2wFl4bSe2r0Zfzn?mode=gi_t

PRINTING AND DISTRIBUTION

Notes Next Door (NND), Delhi is the official printing and distribution partner for Weave Sprint. The soft-copy PDFs and the complete print-ready set are always available at weave.clinic/sprint, should you prefer to print the notes yourself. Weave Sprint is open source and free, and as part of Weave's commitment to residents it will stay that way, forever. We do, however, strongly recommend a printed set from Notes Next Door: in our experience, they have been the most reliable printing partner we have worked with.

WEAVE SPRINT · PAPER IV · VOLUME SIX

Volume Six of the fourth paper is the recent advances chapter: what has actually changed in psychiatric practice and research, and the landmark studies behind each change, read as evidence rather than as an announcement. It is the part of the syllabus that dates fastest, which is why it is bound by itself. One complete set of notes, written to be understood the first time and recalled every time after, with a colour memory code, a clickable contents and a full index.

Built to understand. Kept to revise.

All the very best for your exam.

WEAVE

weave.clinic · wilfred@weave.clinic · niharika@weave.clinic